Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

High-accuracy measurement of external characteristics of new-type power-electronic loads based on a four-virtual-phasor DFT algorithm

Li Chen

chenli271490@163.com

State Grid Electric Power Research Institute (NARI Group Corporation)

Muyong Zhen

State Grid Corporation of China

Feiyang Xu

North China Branch of State Grid Corporation of China

Chi Xu

State Grid Electric Power Research Institute (NARI Group Corporation)

Min Zhang

State Grid Electric Power Research Institute (NARI Group Corporation)

Yu Wang

Southeast University

Article

Keywords: power-electronic load, external characteristic, data-driven modelling, synchrophasor, virtual phasor, discrete Fourier transform, impedance measurement, oscillation risk

Posted Date: November 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11484424/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

High-accuracy measurement of external characteristics of new-type power-electronic loads based on a four-virtual-phasor DFT algorithm

Li Chen[1*]
Muyong Zhen[2]
Feiyang Xu[3]
Chi Xu[1]
Min Zhang[1]
Yu Wang[4]
1. State Grid Electric Power Research Institute (NARI Group Corporation), Nanjing 210003, China
2. National Electric Power Dispatching and Control Center, State Grid Corporation of China, Beijing 100031, China
3. Dispatching and Control Center, North China Branch of State Grid Corporation of China, Beijing 100053, China
4.School of Electrical Engineering, Southeast University, Nanjing 210096, China
* Correspondence: Li Chen, chenli27149@163.com

Abstract

External-characteristic models of new-type power-electronic loads—data-centre rectifier supplies, battery-storage converters and fast chargers—are a prerequisite for the stability analysis of converter-dominated power systems, and both the model-driven and the data-driven modelling routes ultimately rest on the accuracy of the measured voltage and current phasors. This paper addresses the modelling bottleneck from the measurement-algorithm side: a four-virtual-phasor DFT (4VP-DFT) estimator, which analytically cancels the constant and first-order terms of the spectral-leakage-induced dynamic error and shapes an equivalent combined window that suppresses off-bin interferers, is generalized from synchrophasor measurement to the measurement of both faces of the external characteristic of a *2* MVA grid-feeding rectifier-type load. The design properties are established as three propositions with closed-form proofs, and the estimator is benchmarked against plain DFT, virtual-αβ, offset-π/3, interpolated-DFT (ipDFT) and second-order Taylor–Fourier baselines under an extended condition set (frequency offsets to ±10 Hz, harmonic distortion up to 10% THD with 50 dB noise, out-of-band and interharmonic interference, amplitude–phase modulation and magnitude–phase steps) with Monte *C*arlo statistics over 200 randomized realizations. At a −5 Hz offset the maximum phase error is 0.0095°, 37% below the best virtual-phasor baseline, and the 40 ms step response remains within the 70 ms limit; the estimator dominates every single-cycle-window baseline in the single-cycle latency class, whereas the multi-cycle baselines, although more accurate in stationary conditions, require two-to-three-cycle windows and degrade beyond ±5 Hz. In the modelling loop, the static V–I characteristic is tracked with an RMSE of 4.69 kW (62% below the plain DFT), the wideband impedance is

measured in 0.2 s windows with a mean magnitude error of 0.72% (plain DFT: 6.54%), and the impedance-based stability margins reproduce the analytic ranking with a worst-case deviation of 33%. Joint identification of two drifted control parameters from the accurately measured impedance recovers the true values within 2.8% and 10.9% and reduces the model-to-data discrepancy from 55.0% to 8.73%, whereas the same procedure driven by plain-DFT measurements converges to a wrong model. The results establish phasor-measurement accuracy as the shared foundation of data-driven external-characteristic modelling, model validation and measurement-driven model correction for power-electronic loads.

Keywords

power-electronic load; external characteristic; data-driven modelling; synchrophasor; virtual phasor; discrete Fourier transform; impedance measurement; oscillation risk

1. Introduction

The load side of power systems is undergoing a profound transformation towards power-electronic interfacing. The most visible driver is the rapid build-out of AI training data centres: synchronized training iterations alternate between computation-heavy phases, in which every GPU processes local data, and communication-heavy phases, in which all GPUs synchronize, so that large, periodic power swings arise at the iteration frequency, with amplitudes growing with job size [1]. Field-oriented profiling further shows that these periodic stochastic fluctuations differ fundamentally from conventional load ramping: they act as persistent forcing inputs that can interact with and amplify both local and inter-area oscillation modes, and case studies on the WECC 179-bus system indicate that narrower fluctuation bands, larger single-site capacities and geographically dispersed siting can all intensify oscillations across multiple modes [2]. Because data-centre electricity consumption is projected to grow several-fold within this decade [3], load-induced oscillations are becoming a recognized operational risk, and a recent measurement-adequacy assessment warns that oscillations associated with periodic data-centre load profiles span a wide frequency range—from electromechanical to sub-synchronous—while existing phasor measurement units (PMUs) are highly effective for low-frequency electromechanical oscillations but suffer inherent bandwidth and reporting-rate limitations of the phasor-estimation filters, which can significantly attenuate oscillation magnitudes and lead to underestimation of oscillatory behaviour [4]. Effective oscillation governance—damping coordination, protection settings and interconnection screening—therefore hinges critically on an accurate characterization of how these loads actually behave at their grid interface, which in turn is bounded by the accuracy of the phasor measurements from which that characterization is derived.

Accurate load models are the basis of any countermeasure. On the model-driven route, a bottom-up electromagnetic-transient model of an AI data centre has recently been released in PSCAD by ERCOT: it is modular, capturing the rectifier, inverter, bidirectional DC/DC converter, battery storage, variable-frequency-drive cooling and aggregated computing loads, and it implements both switching and average converter models, with the average model substantially improving simulation efficiency while preserving the essential low-frequency dynamics [5]. Such models, however, are assembled from publicly available information and are not validated against actual hardware, so their parameters—controller gains, PLL bandwidths, ride-through settings—remain nominal estimates that drift with firmware updates and commissioning practice. The same limitation is visible in operating experience: in hydro-dominated grids, ultra-low-frequency oscillation induced by daily AGC operation has been observed and had to be suppressed by coordinated battery storage control [6], and sub-synchronous oscillations have been shown to propagate through grids with high wind penetration according to well-defined propagation characteristics [7]. In all of these events the decisive quantities—damping contribution, impedance seen at the interconnection—are properties of the actual installed equipment, not of the nameplate model. A model that cannot be checked, and corrected, against measurements of the real equipment is of limited value for oscillation-risk assessment. The natural remedy is therefore to characterize the load from measurements taken at its own terminals—the data-driven modelling route—whose credibility ultimately rests on a single prerequisite: the measured voltage and current phasors must be accurate enough to be trusted.

What must be modelled is the external characteristic of the load as seen from the grid, which has two complementary faces (Fig. 1). The *static* face is the V–I plane trajectory: a constant-power load draws more current as voltage dips, presenting negative incremental resistance at the operating point, which aggravates voltage stability and fault-response behaviour. The *dynamic* face is the small-signal impedance $Z(f)$ over a wide frequency band, which governs converter–grid interactions; impedance-based criteria, in which the ratio of grid impedance to load impedance must keep a safe distance from the critical point (−1, 0), are now standard practice for oscillation-risk assessment of converter-dominated systems [8–11]. Reliable knowledge of both faces is a prerequisite for online oscillation monitoring, wide-area damping control and planning-stage security checks [12].

Neither face can be computed from nameplate data, because control parameters, operating strategies and ride-through logic are opaque to the system operator; both must therefore be measured from the terminal voltage and current waveforms. This is precisely where data-driven modelling fails today. When the grid frequency deviates from nominal, the single-cycle DFT suffers

spectral leakage that splits into a static component—compensable by frequency tracking [13]—and a dynamic component that depends sinusoidally on the signal initial phase and is difficult to correct in real time [14,15]. Virtual-phasor techniques mitigate the dynamic error by combining the phasor with copies computed on time-shifted windows; representative variants include the virtual-αβ stationary-frame method [35] and the offset-π/3 method [36] (both defined in Table *1).* Reported results, however, exhibit a persistent trade-off that no existing variant escapes: estimators with strong dynamic-error suppression are weak against harmonic interference, and vice versa. Multi-cycle estimators escape this trade-off by widening the main lobe: interpolated DFT (ipDFT) methods *w*ith cosine windows model and cancel the *l*eakage analytically [15,37], and Taylor–Fourier methods fit the time-varying phasor directly [14,38]; both, however, pay with two-to-three-cycle observation spans that delay the report and narrow the usable frequency band, which motivates keeping the estimator within the single-cycle class. For impedance scanning the situation is worse, because the perturbation to be measured is orders of magnitude smaller than the fundamental and is easily buried under leakage from the fundamental itself and from harmonic distortion—exactly the measurement-adequacy concern raised for data-centre oscillation monitoring *[4].* Distorted phasors then propagate through the whole modelling chain: into the identified V–I characteristic and impedance, into the derived stability margins and—if a model-correction loop is built on top—into wrongly identified parameters that corrupt the corrected model. The accuracy of the phasor-measurement layer is thus the common foundation of both the data-driven and the model-driven modelling routes. This observation motivates the research question of this paper: can a single-cycle-window estimator simultaneously achieve second-order suppression of the leakage-induced dynamic error and strong rejection of off-bin harmonic interferers and noise, and does lifting this dual property remove the measurement bottleneck of data-driven external-characteristic modelling?

This paper addresses the modelling bottleneck from the measurement-algorithm side. We generalize a four-virtual-phasor DFT (4VP-DFT) estimator—previously developed by the authors for synchrophasor measurement—to the measurement of both faces of the external characteristic of new-type power-electronic loads, and we embed it in a modelling framework that couples the data-driven route, the model-driven route and a measurement-driven model-correction loop. The main contributions are as follows:

1. A measurement-based modelling framework for power-electronic load external characteristics is used to structure the study, and the phasor-estimation layer is identified as the common bottleneck of data-driven modelling, model-driven validation and *m*easurement-driven *m*odel correction, and an error-propagation analysis, closed form in Eq. *(1),* quantifies ho*w* phasor errors amplify into

stability-margin errors (in proportion to 1/η) and into biased identified parameters; the measured impedance feeds a least-squares identification loop that corrects the model-driven side.

2. A four-virtual-phasor DFT estimator is designed in which two symmetric virtual-phasor pairs with offset angles π/3 and 2π/3 are combined with weights (1/3, 2/9, 1/9). The design rule, derived from a Taylor expansion of the leakage (dynamic-error) coefficient, cancels its constant and first-order terms, so the residual dynamic error scales with the square of the normalized frequency deviation instead of linearly; at the same time, the equivalent combined window provides a third-harmonic suppression coefficient of 1/9—4.5 times smaller than that of the virtual-αβ method—attenuates half-bin-offset interferers to 6% at the perturbation bins and reduces the equivalent noise bandwidth to 19/81 of the rectangular window. To the best of the authors' knowledge, this is the first single-cycle-window estimator to combine second-order dynamic-error cancellation with strong harmonic, interference and noise immunity, breaking the accuracy-versus-robustness trade-off that limits existing virtual-*p*hasor variants.

3. The estimator is validated against the accuracy requirements of synchrophasor standards under an extended condition set—frequency offsets to ±10 Hz, harmonic distortion up to 10% THD with 50 dB noise, out-of-band and interharmonic interference, amplitude–phase modulation and magnitude–phase steps—with Monte Carlo statistics over 200 randomized realizations, and is benchmarked against interpolated-DFT and second-order Taylor–Fourier baselines (Tables *3* and *4)*. At a −5 Hz offset the maximum phase error is 0.0095°, 37% lower than the best reported virtual-phasor variant; the step response is 40 ms, well within the 70 ms limit.

4. The algorithm is applied to a detailed average dq model of a *2* MVA grid-feeding rectifier-type load representing data-centre power supplies, storage converters and fast chargers. With balanced multi-sine injection and fundamental removal, the wideband impedance is measured in 0.2 s windows with a mean magnitude error of 0.72% (plain DFT: 6.54%) and coincides with the ground truth, while the two independent references—the data-driven ground truth and the Jacobian linearization of the average model—agree with each other within 1.0%, establishing the consistency of the data-driven and model-driven routes; the static V–I characteristic under AI-workload-like power fluctuations is tracked with an RMSE of 4.69 kW, a 62% reduction relative to the plain DFT.

5. Closing the modelling loop, two drifted control parameters (PLL bandwidth 15 Hz vs. the true 10 Hz; current-loop bandwidth 300 Hz vs. the true 400 Hz) are jointly identified from the measured impedance by least squares. Joint identification from 4VP-DFT measurements recovers 9.72 Hz and 356 Hz (2.8% and 10.9% errors) and reduces the model-to-data discrepancy from 55.0% to 8.73%, whereas the same procedure driven by plain-DFT measurements misidentifies the PLL bandwidth by 34% and leaves a discrepancy of 14.4%; identification from clean reference

waveforms recovers both parameters exactly, confirming that the parameters are identifiable and that measurement accuracy, not identifiability, decides whether model correction succeeds. A 36-scenario robustness scan over the nominal-parameter plane further shows that the loop retains a median post-correction deviation of 11.0%, degrades gracefully to a detectable saturation when the drift exceeds the retuning authority, and should be gated by the measurement uncertainty.

The remainder of the paper is organized as follows. Section *2* develops the theory: the modelling framework and the role of the measurement layer (2.1) and the 4VP-DFT estimator with its design rules and properties (2.2). Section *3* validates the algorithm: compliance with synchrophasor accuracy standards (3.1)—extended to frequency offsets of ±10 Hz, interharmonic interference and 200-realization Monte Carlo statistics, and benchmarked against interpolated-DFT and Taylor–Fourier baselines—and data-driven measurement of the static external characteristic (3.2). Section *4* presents the simulation analysis on a data-centre-type load: the dynamic external characteristic and its cross-verification against the model-driven route (4.1), measurement-driven model correction with a robustness scan (4.2), and oscillation-risk assessment from the measured characteristic (4.3). Section *5* discusses the engineering implications and limitations, Section *6* concludes the paper, and the Methods detail the load model, the measurement procedures and the simulation platform.

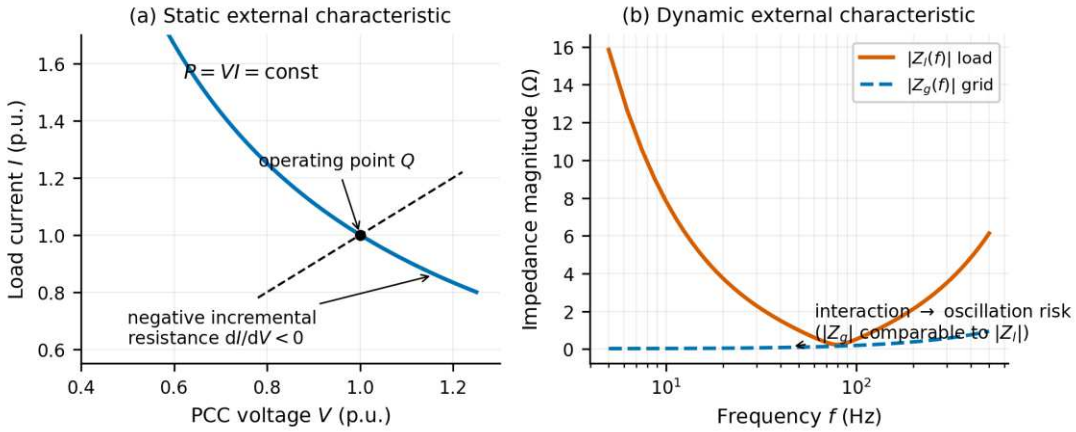


Figure 1. The two faces of the external characteristic of a power-electronic load, the modelling target of this study. (a) Static face: the V–I plane trajectory; constant-power behaviour appears as negative incremental resistance at the operating point Q. (b) Dynamic face: the wideband load impedance $Z_l(f)$ interacting with the grid impedance $Z_g(f)$; when the two become comparable, converter–grid oscillation risk arises.

2. Theoretical analysis

2.1 Modelling framework and the role of the measurement layer

Figure 2 shows the modelling framework underlying this study. The *model-driven route* (upper left) builds on an average converter model of the load—here a grid-feeding rectifier with cascaded DC-voltage, current and PLL controls feeding a constant-power DC side, an abstraction of the rectifier front of data-centre supplies, storage PCS and fast chargers in the spirit of the composite models recently assembled for AI data centres [5]. Linearizing the model at an operating point yields an analytic impedance $Z_m(f; \theta)$ parameterized by the control parameters θ (PLL bandwidths, loop gains, delays). The *data-driven route* (lower left) estimates the same quantities from terminal waveforms: cycle-by-cycle fundamental phasors give the static power trajectory $P(t)$, and frequency-scan processing of the perturbation components gives the wideband impedance $Z(f)$. Both routes feed the characteristic-analysis stage (centre): the negative incremental resistance of the static face, and the impedance-based stability margin η of the dynamic face. Whenever the two routes disagree, the discrepancy can be attributed to drifted model parameters θ, and a least-squares identification from the measured impedance yields corrected parameters $\hat{\theta}$ that update the model-driven side (right)—the measurement-driven model-correction loop.

Every stage of this framework consumes phasors, so a measurement error is not confined to the data-driven route but propagates through the whole chain. For the static face, a phasor error *(δV, δI)* maps directly into an error $\delta P \approx I \cdot \delta V + V \cdot \delta I$ of the tracked power. For the dynamic face, an impedance error δZ at each tone propagates into the stability margin $\eta = \min|1 + Zg/Zl|$ with a sensitivity that grows as the Nyquist curve approaches the critical point: as derived below (Eq. *(1)),* the relative margin error is bounded by $(|Zg/Zl|/\eta)\cdot\delta|Zl|/|Zl|$, so near the oscillation boundary, where the margin is smallest and the assessment matters most, the same relative impedance error is amplified by 1/η. The model-correction stage inherits the same sensitivity: because the identified parameters minimize a cost function built from $\hat{Z}$, a biased impedance displaces the minimum and the correction converges to biased parameters. The hypothesis tested in this paper is therefore that the accuracy of the phasor-estimation layer—not the perturbation hardware—is the binding constraint, because converter terminals exhibit off-nominal frequency, harmonic distortion and noise simultaneously, which is the harshest environment for DFT-based estimators [4,14,15]. Section 2.2 develops the estimator, Section *3* validates it, and Sections *3–4* quantify how its accuracy propagates through the whole modelling loop.

To derive Eq. *(1),* write the loop gain as $L(f) = Zg(f)/Zl(f)$ and let $f^\star = \arg\min_f |1 + L(f)|$, so that $\eta = |1 + L(f^\star)|$. A relative magnitude error $\varepsilon = \delta|Zl|/|Zl|$ of the measured load impedance rescales the loop gain to $(1 + \varepsilon)L(f)$, and the perturbed margin $\eta(\varepsilon) = \min_f |1 + (1 + \varepsilon)L(f)|$ is minimized, to first order, at $f^\star$, giving $d\eta/d\varepsilon|_0 = (|L(f^\star)|^2 + \mathrm{Re}\, L(f^\star))/\eta$. In magnitude, $|d\eta/d\varepsilon| \le |L(f^\star)|$, hence the relative margin error is bounded by

$$\frac{|\delta\eta|}{\eta} \le \left(\frac{|Z_g(f\star)/Z_l(f\star)|}{\eta}\cdot\frac{\delta|Z_l|}{|Z_l|}\right),\ \ \mathrm{f}\star = \arg\min\left|1 + Z_g/Z_l\right| \tag{1}$$

—the relative impedance error amplified by |Zg/Zl|/η: as the Nyquist curve approaches the critical point *(η → 0),* the same measurement error is amplified without bound. This sensitivity is intrinsic to the criterion and independent of the estimator; it quantifies the accuracy budget that the phasor layer must meet.

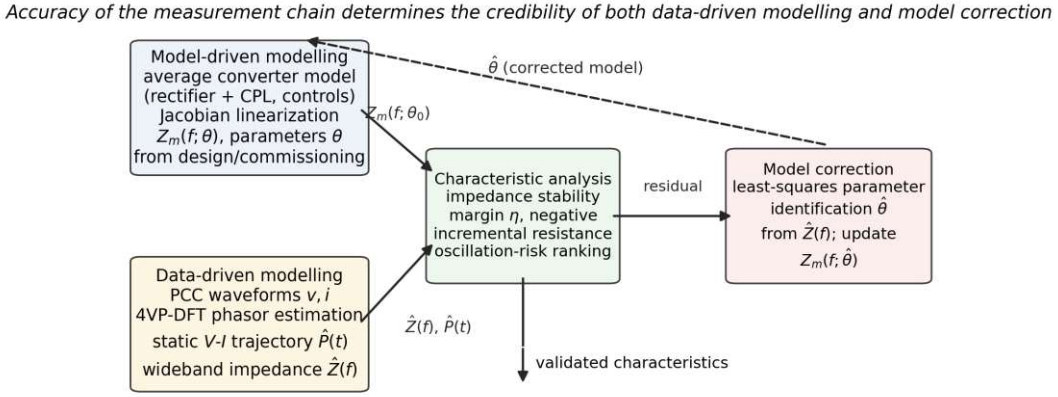


Figure 2. Measurement-based modelling framework for power-electronic load external characteristics. The model-driven route (average converter model, Jacobian linearization $Z_m(f;\theta)$) and the data-driven route (4VP-DFT phasor estimation of the static trajectory $P(t)$ and the wideband impedance $Z(f)$) both feed characteristic analysis; the discrepancy between them drives a least-squares parameter identification that corrects the model. The accuracy of the measurement chain determines the credibility of both routes and of the correction loop.

2.2 Four-virtual-phasor DFT estimator

Signal model and leakage decomposition. Let the sampled signal be

$$x(n) = \sqrt{2}\,A\cos\left(\frac{2\pi(1+r)n}{N} + \varphi\right) + \eta(n) \tag{2}$$

where N is the window length at the nominal frequency f_0, r = (f − f_0)/(f_0/N) is the normalized frequency deviation in bins, φ the initial phase and η(n) disturbance terms. The single-bin DFT phasor at the nominal bin is

$$X(k) = \frac{1}{N}\sum_{n=0}^{N-1} x(n)e^{-\mathrm{j}2\pi kn/N} \tag{3}$$

and, for the signal of Eq. (*2*), it decomposes into the image of the true phasor plus a mirror term [14,15]:

$$\hat{P} = P_0\,W(r) + P_0^{*}\,W(2+r),\ \ W(r) = \frac{\sin(\pi r)}{N\sin\left(\frac{\pi r}{N}\right)}\,e^{j\pi r(N-1)/N} \tag{4}$$

where $P_0 = A\angle\varphi$ and $W(\cdot)$ is the Dirichlet kernel of the rectangular window. The mirror term rotates at twice the signal frequency and is the source of the phase-dependent dynamic error identified in the synchrophasor literature: its magnitude is independent of any static (frequency-tracking) compensation and can only be reduced by combining phasors computed on time-shifted windows.

Virtual phasors. A virtual phasor with offset angle θ is the phasor of a window shifted by $\Delta = N\theta/(2\pi)$ samples, rotated back to the common time reference (Fig. 3):

$$\hat{P}_{\pm\theta} = e^{\pm j\theta}\frac{\sqrt{2}}{N}\sum_{k=0}^{N-1} x(n_0 \mp \Delta + k)e^{-\mathrm{j}2\pi k/N} \quad \Delta = \frac{N\theta}{(2\pi)} \tag{5}$$

Because the shift is realized by re-indexing the same sampled waveform, virtual phasors cost no additional sensing; with a recursive DFT they cost one complex update per new sample per offset. The four-virtual-phasor estimator combines the nominal phasor with two symmetric pairs at $\theta_1 = \pi/3$ and $\theta_2 = 2\pi/3$:

$$\hat{P}_{4\mathrm{vp}} = a\,\hat{P}_0 + b\left(\hat{P}_{+\theta_1} + \hat{P}_{-\theta_1}\right) + c\left(\hat{P}_{+\theta_2} + \hat{P}_{-\theta_2}\right) \tag{6}$$

Design rule. The mirror-term coefficient of the combination is

$$\gamma(r) = a + 2b\cos[(2+\mathrm{r})\theta_1] + 2c\cos[(2+\mathrm{r})\theta_2] \tag{7}$$

Expanding Eq. (*7*) in a Taylor series in r about r = *0,* the constant term is a − b − c (since $\cos 2\theta_1 = \cos(2\pi/3)$ and $\cos 2\theta_2 = \cos(4\pi/3)$ both equal −1/2), and the first-order term is proportional to $-b\theta_1 \sin 2\theta_1 - c\theta_2 \sin 2\theta_2 = (-\sqrt{3}/2)(b\theta_1 - c\theta_2)$. Imposing vanishing constant and first-order terms together with the normalization of the true-phasor gain gives the design equations

$$a - b - c = 0 \tag{8}$$

$$b\theta_1 = c\theta_2 \tag{9}$$

$$a + 2b + 2c = 1 \tag{10}$$

whose rational solution with the smallest observation span is

$$(a, b, c) = \left(\frac{1}{3}, \frac{2}{9}, \frac{1}{9}\right) \quad |A_{\mathrm{H}}(r)| \propto r^2 \tag{11}$$

Consequently the dynamic-error coefficient of the combined estimator satisfies $|AH(r)| = O(r^2)$ (Table *5),* in contrast to O(r) for all single-pair variants. We summarize the analytical properties of the estimator as three propositions, each obtained by direct algebra from Eqs. (6)–(12) and each verified numerically in Section *3.1.* Proposition *1* (dynamic-error cancellation). For a combination of the form *(5)* built from two symmetric pairs (θ_1, θ_2), the constant and first-order terms of the leakage coefficient AH(r) vanish if and only if a = b + c and $b\theta_1 = c\theta_2$; together with the gain normalization

these conditions admit the unique rational solution (10), which has the smallest observation span among two-pair combinations.

$$C(h) = a + 2b\cos[h\theta_1] + 2c\cos[h\theta_2] \quad C(3) = \frac{1}{9} \tag{12}$$

Proposition *2* (on-bin harmonic suppression). The response to an on-bin h-th harmonic is multiplied by Eq. (*12)*, which for the 4VP-DFT gives *C*(*3) = 1*/9. Proposition *3* (off-bin interferer shaping). With respect to an interferer at *δ* bins from the measurement bin, the estimator behaves as a single rectangular window—the classical equivalence underlying window design [37]—shaped by g(δ) = a + 2b·cos(2πδ/6) + 2c·cos(2πδ/3), so that half-bin-offset interferers—the worst case for windowed estimators—are attenuated by |g(0.5)| ≈ 0.83 at the fundamental and by |g(2.5)| ≈ 0.06 at the perturbation bins, versus no shaping for the plain DFT. Taken together, the suppression coefficients of the third harmonic—the strongest odd-harmonic error source for a rectangular window—are *1* for the plain DFT, 1/2 for the virtual-αβ method, 1/3 for the offset-π/3 method and 1/9 for the 4VP-DFT (Table *1).* By the same window equivalence, the noise behaviour is also governed by the shaped window: for additive white noise of variance σ², the phasor-variance ratio to the plain DFT equals the normalized window energy a² + 2b² + 2c² = 19/81 ≈ 0.235, i.e. about *6* dB below the plain DFT, compared with 1/3 for the offset-π/3 method and 3/8 for the virtual-αβ method. All comparison estimators are special cases of the same combination and are defined by their weights and offsets in Table *1.*

Proofs. Proposition *1:* with the mirror-term coefficient γ(r) = a + 2b·cos[2θ$_1$(1−r)] + 2c·cos[2θ$_2$(1−r)], the constant term is γ(0) = a + 2b·cos 2θ$_1$ + 2c·cos 2θ$_2$ = a − b − c and the first-order term is γ′(0) = 4(bθ$_1$·sin 2θ$_1$ + cθ$_2$·sin 2θ$_2$) = 2√3·(bθ$_1$ − cθ$_2$), because sin 2θ$_1$ = √3/2 and sin 2θ$_2$ = −√3/2 for θ$_1$ = π/3 and θ$_2$ = 2π/3. Equations (8)–(10) are therefore exactly the first-order necessary conditions for a minimum of |γ(r)| at r = *0* subject to unit true-phasor gain, and solving them gives (a, b, c) = (1/3, 2/9, 1/9) uniquely; the resulting observation span N(1 + θ$_2$/π) = 5N/3 is the smallest among rational two-pair combinations, because θ$_2$ = 2θ$_1$ is forced by the ratio of Eq. *(9).* Proposition *2:* an on-bin h-th harmonic rotates each shifted-window phasor by hθ, so the combination response is Eq. (12), and with the weights above C(3) = 1/3 − 4/9 + 2/9 = 1/9. Proposition *3:* by the discrete-time Fourier shift theorem, shifting a window by *Δ* samples multiplies its transform by e^(−j2πδΔ/N); a symmetric pair therefore contributes 2b·cos(2πδΔ$_1$/N) + 2c·cos(2πδΔ$_2$/N) = 2b·cos(2πδ/6) + 2c·cos(2πδ/3), since Δ$_1$ = N/6 and Δ$_2$ = N/3, which is the shaped-window equivalence of Proposition *3;* evaluation at *δ* = 0.5 and *δ* = 2.5 gives |g(0.5)| ≈ 0.83 and |g(2.5)| ≈ 0.06.

Reporting and transition. Phasors are reported every *N* samples in the RMS convention of [16]. During large steps the five windows of the 4VP-DFT straddle the discontinuity; a transitional logic

that detects a report-to-report jump beyond 5% and temporarily switches to the shorter offset-π/3 span keeps the dynamic overshoot bounded without affecting steady-state accuracy (Section 3.1). Formally, let J(k) = |P̂(k) − P̂(k−1)|/|P̂(k−1)| denote the relative jump between consecutive reports; the estimator switches to the offset-π/3 span whenever J(k) exceeds 5% and returns to the full 4VP span only after three consecutive reports with *J* below 2%, a hysteresis band that prevents chattering at the switching boundary. Remark *1* (computational cost). The estimator evaluates five N-point DFTs per report; with the recursive DFT update each new sample requires five complex rotations—five times the per-sample cost of the plain DFT and negligible compared with the integration cost of the plant model (Methods).

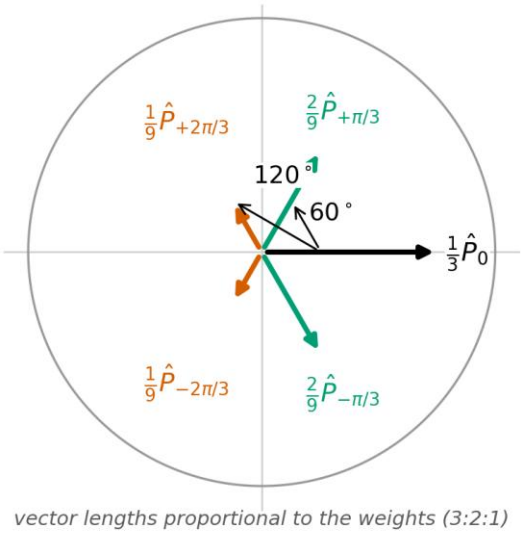


Figure *3.* Construction of the four-virtual-phasor estimator: the nominal phasor (weight 1/3) plus two symmetric pairs computed on windows shifted by ±π/3 (weight 2/9 each) and ±2π/3 (weight 1/9 each), rotated to the common time reference; the combination cancels the constant and first-order leakage terms (Eqs. *(8)*–(10)).

Table *1.* Definition and properties of the six phasor estimators compared in this study *(N* = samples per cycle at 50 Hz; r = normalized frequency deviation in bins; for Hann-ipDFT and TF-2 the leakage is modelled rather than combined, so the weight convention and C(3) do not apply in the same sense).

Estimator	**Weights (a; b, θ_1; c, θ_2)**	**Observation span**	**Dynamic-error scaling**	**3rd-harmonic suppression *C*(3)**
Plain DFT	1; —; —	N (20 ms)	*O*(*r*)	1
Virtual-αβ	1/2; 1/4, π/2; —	4N/3 (26.7 ms)	*O*(r^2)	1/2
Offset-π/3	1/3; 1/3, π/3; —	4N/3 (26.7 ms)	*O*(*r*)	1/3

Proposed 4VP-DFT	1/3; 2/9, π/3; 1/9, 2π/3	5N/3 (33.3 ms)	$O(r^2)$	1/9
Hann-ipDFT [15,37]	— (two-cycle periodic-Hann window; leakage model fit)	2N (40 ms)	model-based (exact static)	1 (3rd on-bin for 2N)
TF-2 [14,38]	— (second-order Taylor basis)	3N (60 ms)	$O(r^3)$, in-band	1 (3rd on-bin for 3N)

3. Algorithm validation

3.1 Compliance with synchrophasor accuracy standards

Before the estimator is embedded in the modelling loop, its own accuracy must be certified, because every downstream quantity inherits its error. The estimator was therefore characterized with the test signals of the synchrophasor testing standard adopted by Chinese grid operators [17]; a sampling rate of 9.6 kHz *(N* = 192 samples per cycle at 50 Hz) was used throughout, and errors are reported as the maximum phase error over an initial-phase sweep and as the total vector error of Eq. (13), evaluated against the theoretical phasor of the test signal, in the RMS convention of IEEE Std C37.118.1-2011 [16]. The Chinese standard [17] specifies the same steady-state and dynamic test signals as IEEE Std C37.118.1-2011—frequency offsets, harmonic-plus-noise mixtures, out-of-band interference, amplitude–phase modulation and magnitude–phase steps—so the compliance results translate directly between the two frameworks. Two extensions beyond the standard condition set are introduced: the frequency range is widened to ±10 Hz, harmonic distortion is raised to 10% THD, interharmonic and dual-tone out-of-band interferers are injected at 10% amplitude, and all noisy conditions are repeated over 200 Monte Carlo realizations with randomized harmonic phases and additive white noise, with errors reported as the mean ± standard deviation across realizations and 95% confidence intervals (Table *4).* To position the estimator against the state of the art, two multi-cycle baselines are benchmarked alongside the single-cycle virtual-phasor family of Table *1:* an interpolated DFT with a two-cycle periodic-Hann window (Hann-ipDFT) [15,37], which fits the image-corrected leakage model of the windowed tone, and a second-order Taylor–Fourier estimator over three cycles (TF-2) [14,38], which solves a normalized second-order Taylor–Fourier least-squares problem. Response times follow the time-tag convention of the standards, in which the estimate is referred to the newest sample of its observation window.

$$\mathrm{TVE} = \frac{\sqrt{(X_\mathrm{r} - \widehat{X_\mathrm{r}})^2 + (X_\mathrm{i} - \widehat{X_\mathrm{i}})^2}}{\sqrt{{X_\mathrm{r}}^2 + {X_\mathrm{i}}^2}} \times 100\% \tag{13}$$

evaluated against the theoretical phasor of the test signal, in the RMS convention of IEEE Std C37.118.1-2011 [16].

Frequency offset. Figure *4* shows the maximum phase error as the signal frequency moves from −5 Hz to +5 Hz around nominal. The proposed 4VP-DFT attains 0.0095° at −5 Hz and 0.0076° at +5 Hz, compared with 0.0151°/0.0136° for the virtual-αβ method and 0.152°/0.129° for the offset-π/3 method; the plain DFT reaches 2.44°. The 4VP-DFT therefore improves the best reported virtual-phasor variant by 37% at −5 Hz, and its error grows quadratically with the normalized deviation (Eq. (*11)* and Table *5*), which explains the flattening around nominal frequency. The corresponding TVEs at −5 Hz are 2.37% for the 4VP-DFT, 2.27% for the virtual-αβ method and 2.32% for the offset-π/3 method (Table *2)*. Although these values are nearly identical across the three estimators, their phase errors differ by more than an order of magnitude, which shows that the residual TVE is a common-mode floor set by the rectangular window's magnitude response and the reference-phasor convention under off-nominal frequency, rather than by the leakage-induced dynamic error targeted in this paper; Extending the sweep to ±10 Hz (Table *3)* preserves the ordering: at +10 Hz the phase errors of the offset-π/3, virtual-αβ and 4VP-DFT estimators are 0.60°, 0.13° and 0.07°, growing from ±5 Hz by a factor of 4–8, consistent with the quadratic law up to higher-order terms, whereas the plain DFT grows by the linear-law factor of two. The Hann-ipDFT baseline maintains machine-level accuracy at every static offset, as expected from its explicit leakage model; the TF-2 baseline retains sub-degree phase accuracy within ±5 Hz but degrades to the common window floor (9.5% TVE) at +10 Hz, the edge of its three-cycle design band. We return to this point in the Discussion.

Harmonics, noise and out-of-band interference. Under a 5.5% THD mixture (2nd–25th harmonics) plus 50 dB white noise, the maximum phase error of the 4VP-DFT is 0.0439°, 17% and 25% lower than the offset-π/3 (0.0531°) and virtual-αβ (0.0586°) methods (Table *2,* Fig. *5)*. Under out-of-band interference, the 4VP-DFT achieves 0.1766° and a TVE of 0.666%, improving the offset-π/3 method by 26% and 38% and the virtual-αβ method by 41% and 52%. Under simultaneous amplitude–phase modulation all single-cycle-window estimators exhibit similar errors (≈4.3°, TVE ≈6.3%), because the upper modulation sideband coincides with the measurement frequency and cannot be rejected by any one-cycle window; this is an inherent limitation of the window length rather than of the estimator design. Monte Carlo repetition (200 realizations with randomized harmonic phases; Table *4)* confirms that the harmonic-condition results are robust rather than incidental: at 10% THD with 50 dB noise the mean maximum phase errors are 0.011° ± 0.006° (plain DFT), 0.010° ± 0.006° (4VP-DFT), 0.125° ± 0.008° (Hann-ipDFT) and 0.333° ± 0.165° (TF-2), with 95% confidence half-widths of at most 0.023°. Under dual out-of-band tones at 125 Hz and 175 Hz (10% each, 50 dB noise) the 4VP-DFT attains 0.14° ± 0.05° and 0.26% ± 0.09% TVE on average—10.0 and 9.9 times lower than the plain DFT—while the Hann-ipDFT is best in this

condition (0.03° ± 0.01°) thanks to its −31 dB sidelobes. Under a 10% interharmonic at f_0 + 125.5 Hz the same ordering holds (4VP-DFT 0.06°/0.10%; Hann-ipDFT 0.01°/0.02%).

Step response. For a simultaneous 10% magnitude and 10° phase step, the 4VP-DFT settles within the 1% TVE band in 40 ms and the plain DFT in 20 ms (Fig. *6),* both satisfying the 70 ms limit of [16]. The two multi-cycle baselines also satisfy the 70 ms limit (Hann-ipDFT 40 ms; TF-2 40 ms), the latter because its end-point evaluation offsets part of its 60 ms span; their straddle overshoots (18.9%) are comparable to those of the single-cycle family (17.9–18.9%), so the observation span rather than the window design dominates this test. The transitional switching to a shorter-window estimator during the step keeps the overshoot bounded without degrading steady-state accuracy.

Analytical dynamic-error coefficient. Table *5* lists the analytical magnitude |AH(r)| of the dynamic-error coefficient at r = 0.1. Single-pair virtual-phasor variants scale linearly with r (offset-π/6: 9.2×10^{-2}; offset-π/4: 7.9×10^{-2}; offset-π/3: 5.9×10^{-2}), whereas the two-pair virtual-αβ and 4VP-DFT estimators scale quadratically; the 4VP-DFT coefficient (3.4×10^{-3}) is 44% lower than that of the virtual-αβ method (6.2×10^{-3}), confirming Eqs. (*8)–(11*) experimentally.

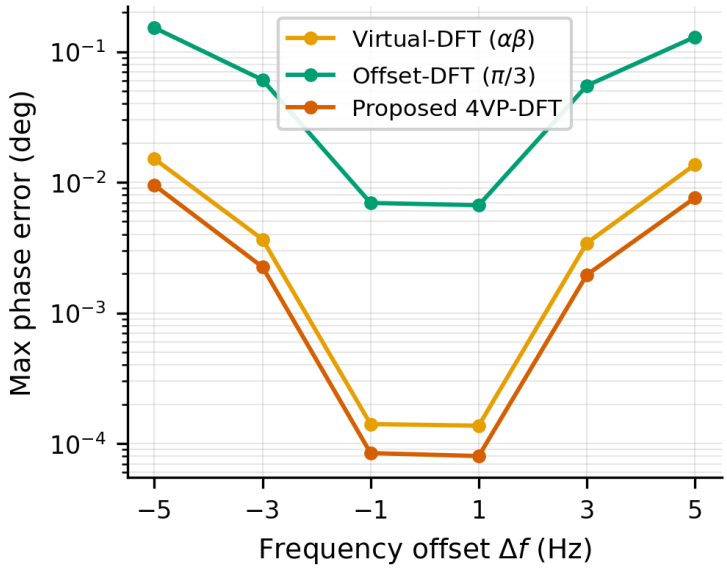


Figure *4.* Maximum phase error over an initial-phase sweep versus frequency offset *(N* = 192). The 4VP-DFT error is quadratic in the normalized deviation (Eq. (11)) and remains below 0.012° at ±5 Hz and 0.11° at −10 Hz.

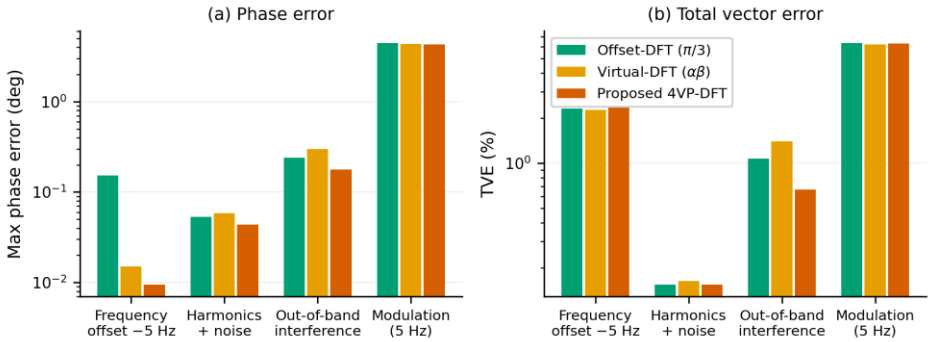


Figure *5.* Compliance results under the standard test conditions: (a) maximum phase error; (b) total vector error (log scale). Under modulation all estimators of the single-cycle family are equally limited by the window length (Table *3)*.

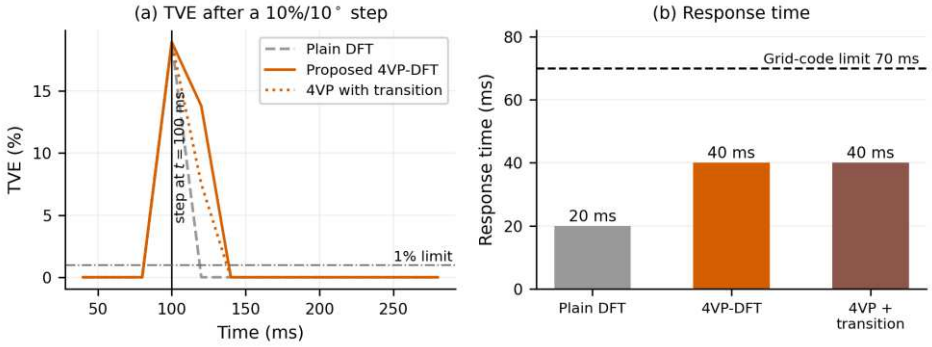


Figure 6. Step-response test (10% magnitude and 10° phase step at t = 100 ms). (a) TVE of the estimators; the 1% limit defines the response time. (b) Response times versus the 70 ms grid-code limit.

Table 2. Compliance results: maximum phase error (deg) and TVE (%) under the test conditions of the Methods (N = 192).

Condition	Offset-π/3: phase/TVE	Virtual-αβ: phase/TVE	4VP-DFT: phase/TVE
Frequency offset −5 Hz	0.1523 / 2.323	0.0151 / 2.274	0.0095 / 2.374
Harmonics (5.5% THD) + 50 dB noise	0.0531 / 0.155	0.0586 / 0.164	0.0439 / 0.155
Out-of-band interference	0.2384 / 1.073	0.2981 / 1.398	0.1766 / 0.666
Amplitude–phase modulation	4.4655 / 6.372	4.3164 / 6.224	4.3157 / 6.289
Step response time (limit 70 ms)	—	— (plain DFT: 20 ms)	40 ms

Beyond the standard condition set, the benchmark is extended in two directions: the condition envelope and the estimator family. Table *3* reports the cross-family comparison under frequency offsets to ±10 Hz, amplitude–phase modulation and the magnitude–phase step; Table *4* reports Monte Carlo statistics under harmonic-plus-noise, out-of-band and interharmonic conditions with

200 realizations per condition. The benchmark leads to three conclusions. First, within the single-cycle latency class the 4VP-DFT strictly dominates: at −10 Hz its phase error is 0.106°, 7.8 times lower than the offset-π/3 method and 1.5 times lower than the virtual-αβ method, and the gap widens quadratically with the offset. Second, the multi-cycle baselines confirm the latency–accuracy trade-off rather than overturn it: the Hann-ipDFT is the most accurate estimator in every stationary condition—its leakage model is exact for a single tone—but its two-cycle window doubles the reporting latency and its modulation error is twice that of the 4VP-DFT; the TF-2 estimator is the only baseline that partially rejects near-fundamental interferers, but its three-cycle window triples the latency, its accuracy collapses at frequency offsets beyond ±5 Hz, and it is the most sensitive of the six estimators to wideband harmonic distortion. Third, under wideband harmonic-plus-noise distortion the single-cycle family is insensitive to the distortion spectrum—the 4VP-DFT, virtual-αβ and offset-π/3 estimators are statistically indistinguishable at the 10% THD level—because their rectangular-window noise floors dominate; the estimator choice is therefore decided by the frequency-offset, interference and latency conditions, exactly the regime of impedance scanning near the fundamental.

Table *3.* Cross-family comparison under the extended test conditions *(N* = 192; maximum phase error in deg and TVE in *%,* maxima over a 289-point initial-phase grid; step: response time with the maximum TVE within 90 ms after the step in parentheses; modulation: ka = 0.1, kp = 0.1 rad, fm = *5* Hz; the 70 ms limit applies to the step test).

Test condition	**Plain DFT**	**Offset-π/3**	**Virtual-αβ**	**4VP-DFT (proposed)**	**Hann-ipDFT**	**TF-2**
Frequency offset −5 Hz	3.017 / 6.82	0.188 / 2.32	0.019 / 2.27	0.012 / 2.37	0.000 / 0.00	0.098 / 0.79
Frequency offset +5 Hz	2.730 / 6.32	0.160 / 2.27	0.017 / 2.27	0.009 / 2.37	0.000 / 0.00	0.102 / 0.83
Frequency offset −10 Hz	6.380 / 16.85	0.823 / 9.14	0.160 / 8.99	0.106 / 9.33	0.000 / 0.00	0.750 / 9.10
Frequency offset +10 Hz	5.217 / 14.96	0.596 / 8.77	0.131 / 8.95	0.068 / 9.26	0.000 / 0.00	0.790 / 9.47
Amplitude–phase modulation	1.72 / 3.65	1.75 / 3.16	1.73 / 3.18	1.73 / 3.18	3.37 / 6.16	3.49 / 6.39
Step response (limit 70 ms)	20 ms (18.9)	40 ms (18.5)	40 ms (18.0)	40 ms (17.9)	40 ms (18.9)	40 ms (18.9)

Table *4.* Monte Carlo statistics of the maximum phase error *(°)* and TVE *(%)* over an eight-point initial-phase grid: mean ± standard deviation across 200 realizations with randomized harmonic phases and 50 dB white noise (95% confidence half-widths do not exceed 0.46° and 0.80% TVE across all entries).

Condition (50 dB white noise throughout)	**Plain DFT**	**Offset-π/3**	**Virtual-αβ**	**4VP-DFT (proposed)**	**Hann-ipDFT**	**TF-2**
Harmonics, 5.5% THD	0.011±0.006 / 0.021±0.010	0.010±0.005 / 0.019±0.009	0.010±0.005 / 0.018±0.009	0.010±0.005 / 0.018±0.009	0.070±0.009 / 0.128±0.015	0.196±0.086 / 0.350±0.153
Harmonics, 10% THD	0.011±0.006 / 0.020±0.011	0.010±0.006 / 0.018±0.011	0.010±0.006 / 0.018±0.010	0.010±0.006 / 0.018±0.010	0.125±0.008 / 0.230±0.014	0.333±0.165 / 0.594±0.295
OOB dual tones, f_0+75 Hz and f_0+125 Hz, 10% each	1.416±0.679 / 2.515±1.189	0.441±0.183 / 0.789±0.324	0.536±0.261 / 0.958±0.462	0.142±0.049 / 0.255±0.086	0.032±0.009 / 0.059±0.016	0.306±0.163 / 0.551±0.294
Interharmonic, f_0+125.5 Hz, 10%	0.767±0.279 / 1.361±0.484	0.207±0.093 / 0.365±0.161	0.343±0.077 / 0.613±0.136	0.058±0.028 / 0.103±0.050	0.011±0.006 / 0.020±0.010	0.161±0.071 / 0.293±0.132

Table *5.* Analytical dynamic-error coefficient |AH(r)| at r = 0.1 and its scaling with r.

Estimator	**$\lvert A_H(0.1)\rvert$**	**Scaling**
Offset-π/6	9.20×10^{-2}	$\propto r$
Offset-π/4	7.85×10^{-2}	$\propto r$
Offset-π/3	5.85×10^{-2}	$\propto r$
Virtual-αβ	6.16×10^{-3}	$\propto r^2$
Proposed 4VP-DFT	3.43×10^{-3}	$\propto r^2$

3.2 Data-driven modelling of the static external characteristic

The static face of the external characteristic was measured on the average model of the 2 MVA grid-feeding rectifier-type load (Methods) whose DC side draws a constant power $P(t)$ following a three-tone fluctuation pattern emulating the periodic stochastic power swings of AI training workloads [1,2,5]. The fundamental phasors of the terminal voltage and current were estimated each cycle with the four estimators of Table 1 and combined as

$$\hat{P}(t) \;=\; 3\,Re\{\hat{V}\;\hat{I}^{*}\} \tag{14}$$

Figure 7(a) shows that all virtual-phasor estimators track the true power from the model states closely, while the plain DFT exhibits a visible ripple and bias because the grid frequency in the test was 49.5 Hz, i.e. off-nominal for the one-cycle window. The tracking RMSE over 8 s is 4.67–4.69 kW for the three virtual-phasor estimators, against 12.33 kW for the plain DFT—a 62% reduction (Fig. 7(b)); relative to the 1.6 MW operating power this is 0.29% versus 0.77%. In the V–I plane the voltage and current perturbations have opposite signs, making the negative incremental resistance of the constant-power behaviour directly visible: this is exactly the signature that voltage-stability studies and protection settings near converter-dominated substations require, and it is recovered distortion-free only when the phasor layer is accurate at off-nominal frequency.

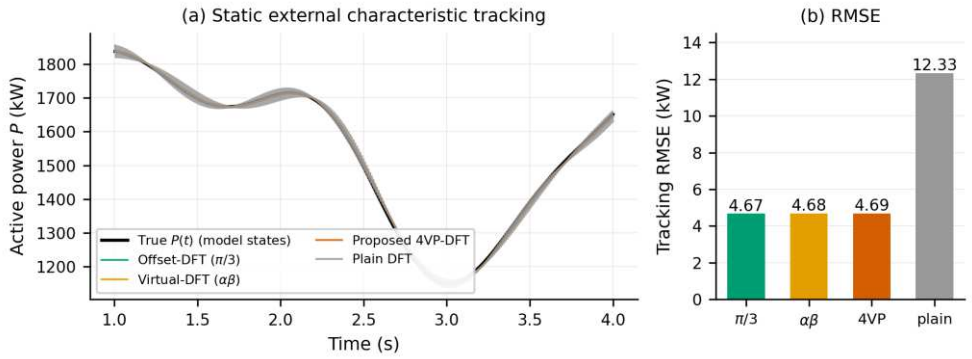


Figure 7. Data-driven measurement of the static external characteristic under AI-workload-like power fluctuations (grid frequency 49.5 Hz): (a) measured active power versus model truth; (b) tracking RMSE of the four estimators.

4. Simulation analysis

4.1 Data-driven modelling of the dynamic external characteristic and comparison with the model-driven route

The dynamic face was obtained by frequency scanning (Fig. 8): a balanced positive-sequence perturbation with 16 tones from 25 Hz to 475 Hz (2% per-tone voltage amplitude) was injected at the point of common coupling while the grid frequency remained 49.5 Hz. Because the perturbation is two orders of magnitude smaller than the fundamental, the known-frequency (49.5 Hz) fundamental was first removed from each analysis window by least-squares fitting of a three-parameter (DC, cosine, sine) model over a 1.2 s span surrounding the window:

$$\sum_k |x(k) - A_0 - A_c \cos(2\pi f_1 t_k) - A_s \sin(2\pi f_1 t_k)|^2 \quad (15)$$

$$\underset{A_0, A_c, A_s}{}$$

which keeps the fit basis orthogonal to the scan tones: with the *5* Hz tone spacing on the *0.2* s grid and the 1.2 s fit span, the correlation between the fitted fundamental and a scan tone is bounded

by |sinc(πΔf_k·T_LS)|, below 0.01% for every grid tone, and the ten-window averaging suppresses the residual further. The removal is essential because the *563 V* fundamental is two orders of magnitude larger than the 11.3 *V* per-tone perturbation (a ≈34 dB ratio) and therefore exceeds the dynamic range that combined-window shaping alone could provide. Residual fundamental, harmonic distortion (2% 5th, 1.2% 7th, placed at half-bin offsets as a stress test) and 50 dB noise remained in the analysis signal. The impedance at each tone was then computed from the phasor ratio

$$\hat{Z}(f) = \frac{\hat{V}(f)}{\hat{I}(f)} \tag{16}$$

in 0.2 s sliding windows and averaged over ten windows.

Figure 9 compares the measured impedance with two independent references: the 10-s Hann-windowed DFT of the clean waveforms (ground truth) and the Jacobian linearization of the average model evaluated at the 49.5 Hz operating point—the model-driven route of Fig. 2. The two references agree to within 1% at every tone, which confirms that the measured impedance is physically consistent with the model: the data-driven and model-driven routes are consistent with each other *when the measurement layer is accurate*. The 0.2 s 4VP-DFT estimates coincide with the ground truth across the whole band, whereas the plain DFT deviates dramatically near 55 Hz, where the impedance peaks and the leakage from the nearby fundamental is strongest.

Quantitatively (Table *6*, Fig. 10), the mean impedance-magnitude error over the 16 tones is 0.72% for the 4VP-DFT, 0.96% for the offset-π/3 method, 4.25% for the virtual-αβ method and 6.54% for the plain DFT; the maximum errors are 1.70%, 2.53%, 40.8% and 70.5%. The phase errors follow the same ordering (0.71°, 0.74°, 1.12°, 2.94° on average). The superiority of the 4VP-DFT stems from the shaping of its equivalent combined window: the combination a + 2b·cos(2πδ/6) + 2c·cos(2πδ/3) of the five shifted rectangular windows attenuates off-bin interferers by up to 25 dB more than the plain window at the half-bin offsets of the injected harmonics, while leaving the on-bin tones unbiased. This is the decisive property for data-driven dynamic modelling: with a plain DFT the impedance near the fundamental—the very region where converter–grid interaction is strongest—is not measurable at all.

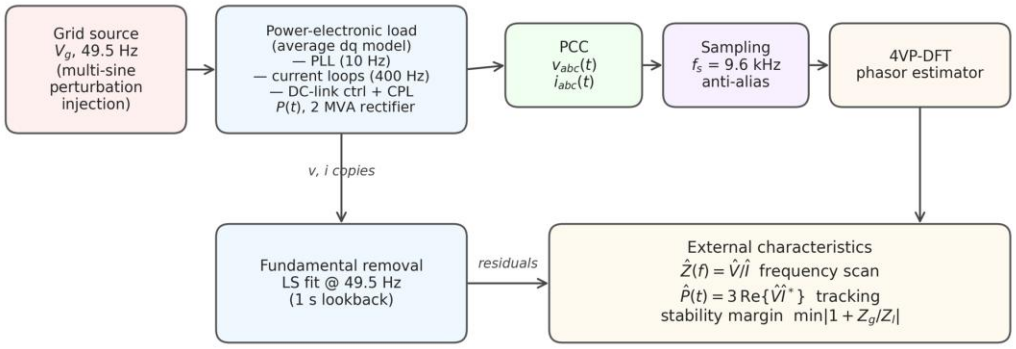


Figure *8.* Measurement chain of the data-driven route: grid-feeding converter load with multi-sine injection, PCC voltage/current acquisition at 9.6 kHz, known-frequency fundamental removal (Eq. (*15)*), 4VP-DFT phasor estimation, and computation of the static and dynamic external characteristics.

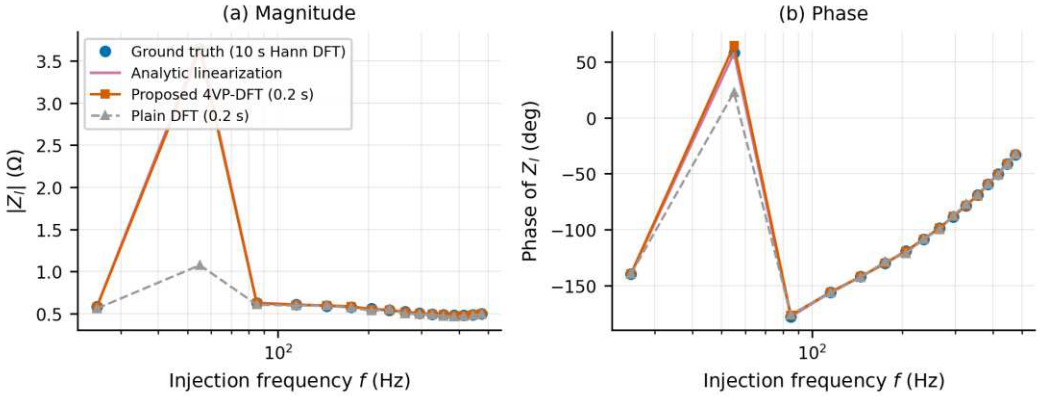


Figure 9. Measured dynamic external characteristic (wideband impedance) of the 2 MVA rectifier-type load: (a) magnitude; (b) phase. Ground truth: 10 s Hann-windowed DFT of the clean record; model-driven reference: Jacobian linearization of the average model (the two references agree within 1.0%). The 0.2 s 4VP-DFT estimates coincide with the truth, while the plain DFT fails near the 55 Hz impedance peak.

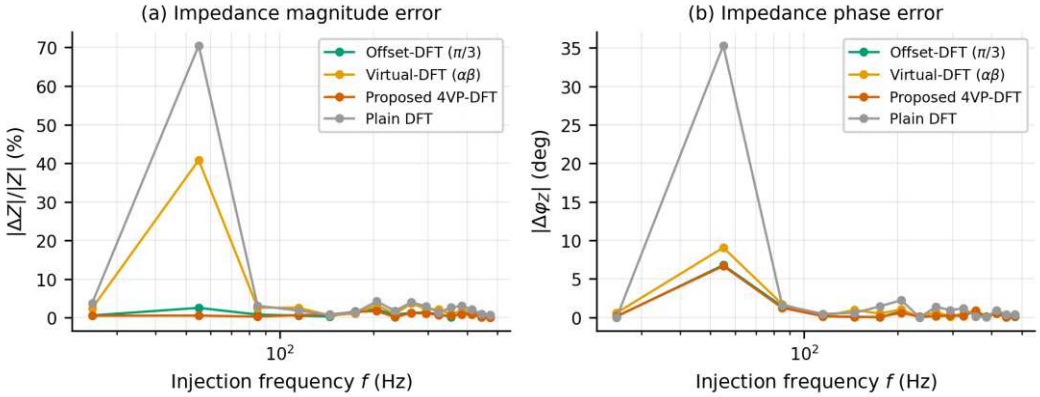


Figure 10. Impedance estimation errors per tone (log scale): (a) magnitude; (b) phase. The combined-window shaping of the 4VP-DFT suppresses the half-bin-offset harmonic interferers that dominate the plain-DFT and virtual-αβ errors.

Table *6*. Impedance-scan errors of the four estimators (16 tones, 0.2 s windows, harmonics + noise stress).

Estimator	Mean $\|\Delta Z\|/\|Z\|$	Max $\|\Delta Z\|/\|Z\|$	Mean $\|\Delta\varphi_Z\|$ (deg)	Max $\|\Delta\varphi_Z\|$ (deg)
Plain DFT	6.54%	70.51%	2.94	35.30
Virtual-αβ	4.25%	40.80%	1.12	9.07
Offset-π/3	0.96%	2.53%	0.74	6.79
Proposed 4VP-DFT	0.72%	1.70%	0.71	6.68

4.2 Measurement-driven model correction: accuracy decides success

We now close the loop of Fig. 2. In service, model parameters drift from their nominal values—controller gains are retuned, firmware is updated, commissioning values differ from design values—so the model-driven impedance $Z_m(f;\ \theta_0)$ disagrees with the measured $Z(f)$ even though the model structure is correct. To represent this realistically, the load model was run with a true PLL bandwidth of 10 Hz and a true current-loop bandwidth of 400 Hz, while the model-driven side used the nominal (drifted) values θ_0 = (15 Hz, 300 Hz). Figure 11 shows the resulting discrepancy: the nominal model deviates from the measured impedance by 55.0% on average and by up to 507% at individual tones—far too large for any stability conclusion. A joint least-squares identification of both bandwidths from the measured impedance,

$$\hat{\beta} = \sum_k |Z_{\mathrm{meas}}(f_{\mathrm{k}}) - Z_{\mathrm{model}}(f_{\mathrm{k}};\ \beta)|^2 \quad (17)$$

was then performed using the measured tone impedances as Zmeas and the parameterized analytic linearization Zmodel(fk; *β)* (Methods) as the model, with *β* = *(fpll, fcur)* the parameter vector; the search is a coarse two-dimensional grid followed by coordinate-descent refinement; the coordinate descent evaluates the one-dimensional least-squares residual afresh at every trial point and terminates when the halving step falls below the refinement resolution (0.02 Hz for the PLL bandwidth, 0.8 Hz for the current-loop bandwidth), so no gradient tolerance is involved (Methods).

The outcome depends entirely on which estimator produced Zmeas (Table *7*). Joint identification from the 4VP-DFT impedance converges to (9.72 Hz, 356 Hz)—parameter errors of 2.8% and 10.9% relative to the true (10 Hz, 400 Hz)—and the corrected model reproduces the measured impedance with a mean deviation of 8.73% and a maximum of 14.8% (Fig. 11). The same procedure driven by the plain-DFT impedance converges to (6.56 Hz, 456 Hz): the PLL bandwidth is misidentified by 34%, and even after correction the model still deviates by 14.4% on average

(69.4% at the worst tone)—the correction loop has converged to a wrong model. A control run using the clean 10 s reference waveforms instead of the 0.2 s estimates recovers both parameters exactly (10.00 Hz, 400.00 Hz), which confirms that the two bandwidths are jointly identifiable from the 16-tone scan and that the limiting factor is the measurement accuracy, not the information content of the data. This control experiment is the methodological cornerstone of the correction loop: because the clean reference recovers both parameters exactly, the residual identification error of the 0.2 s window estimates can be attributed unambiguously to measurement noise rather than to weak identifiability. Measurement accuracy therefore decides whether model correction succeeds—a distorted measurement does not merely degrade the data-driven route, it actively corrupts the model-driven route when the two are coupled.

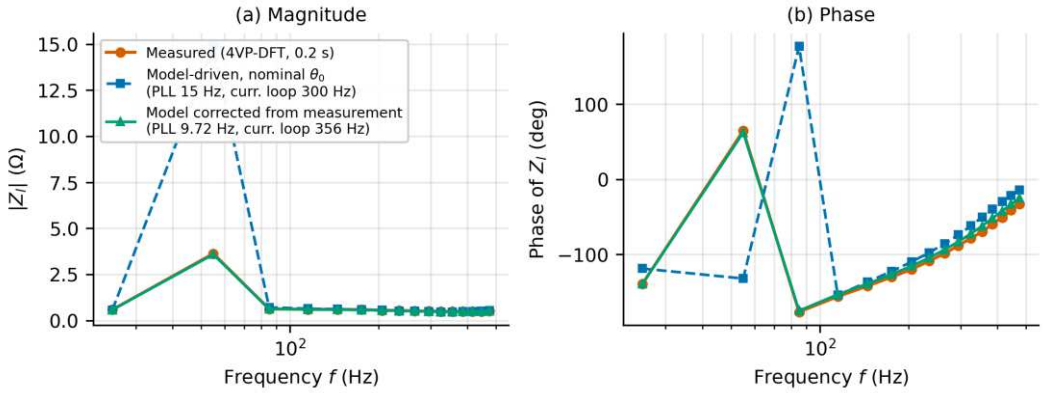


Figure 11. Measurement-driven model correction. (a) Magnitude and (b) phase of the load impedance: measured with the 4VP-DFT (0.2 s windows), predicted by the model-driven route with the nominal (drifted) parameters (PLL 15 Hz, current loop 300 Hz), and predicted after joint identification (PLL 9.72 Hz, current loop 356 Hz; true values 10 Hz and 400 Hz).

Table 7. Joint measurement-driven identification and correction of drifted PLL and current-loop bandwidths (true values 10 Hz and 400 Hz; nominal model values 15 Hz and 300 Hz).

Quantity	4VP-DFT measurement	Plain-DFT measurement
Identified (PLL, current loop)	9.72 Hz, 356 Hz (err. 2.8% / 10.9%)	6.56 Hz, 456 Hz (err. 34% / 14%)
Model-to-data deviation before correction (mean / max)	55.0% / 507%	55.0% / 507%
Deviation after correction (mean / max)	8.73% / 14.8%	14.4% / 69.4%
Deviation with clean 10 s reference (identifiability check)	1.76% / 12.0%	1.76% / 12.0%

The correction loop was further stress-tested over a 6 × 6 grid of drift scenarios spanning nominal PLL bandwidths $f_{pll,0} \in \{6, 8, 10, 12, 15, 18\}$ Hz and nominal current-loop bandwidths $f_{cur,0} \in \{250,$

300, 350, 400, 450, 500} Hz—single-parameter drifts, joint drifts, and the no-drift cell—with the plant fixed at (10 Hz, 400 Hz). To emulate the limited authority of an online retuning procedure, identification was restricted to a ±50% window around each nominal pair using a three-stage nested grid search inside the window; this local procedure is deliberately different from the global search described in the Methods and probes whether the correction concept remains usable when the model cannot be freely re-parameterized. Three findings emerge (Fig. 12). First, robustness: across all 36 combinations the corrected model deviates from the true plant impedance by a median of 11.0% (mean 12.3%, worst 19.8%) when driven by 4VP-DFT measurements, against 14.8% (mean 15.8%, worst 32.0%) for plain-DFT, and the identified PLL bandwidth stays within 4.6% of the true value on average (29% for plain-DFT); for the 35 genuinely drifted cells the correction reduces the pre-correction deviation by a median factor of 2.7, up to 10.2. Second, the applicability boundary: when the nominal PLL bandwidth underestimates the true value beyond the retuning authority (the $f_p$ll,0 = 6 Hz column, whose window is capped at 9 Hz), the identified bandwidth saturates at the window edge and the residual deviation stays at 19.3–19.8%, visible as the dark saturated column of Fig. 12(a). The loop thus fails gracefully—rather than issuing a meaningless update it signals that the plant has left the validity domain of the model, where re-modelling rather than re-tuning is required. Third, a gating rule: in the no-drift cell the nominal model already matches the measurement within 1.76% on average, below the measurement's own worst-tone uncertainty of 12.0%, yet blind correction still moves the parameters and degrades the model-to-true deviation to 10.7%. Correction should therefore be triggered only when the pre-correction residual exceeds the measurement uncertainty bound; otherwise the nominal model is retained. The correction loop is thus accurate where it should act, honest about where it must not, and its value is greatest exactly when the measurement chain is accurate enough to expose the drift.

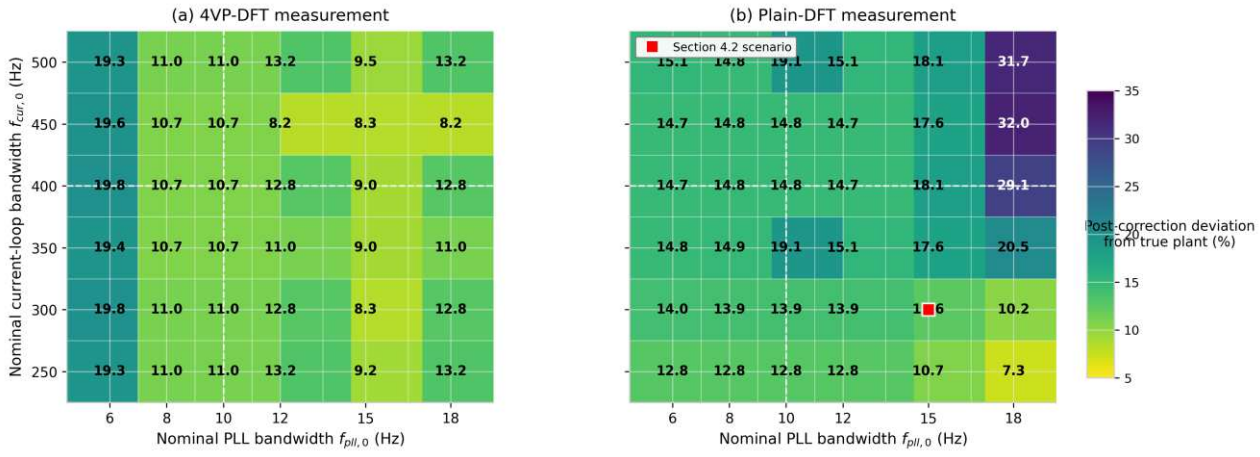


Figure 12. Robustness of the measurement-driven correction loop over 36 drift scenarios. Post-correction deviation of the corrected model from the true plant impedance (%) over the plane of nominal parameters,

with identification restricted to a ±50% retuning window around each nominal pair: (a) identification from 4VP-DFT measurements; (b) identification from plain-DFT measurements. White dashed lines: true values (10 Hz, 400 Hz); red square: the Section 4.2 scenario (15 Hz, 300 Hz). The saturated column at $f_p$ll,0 = 6 Hz marks the applicable boundary of the correction authority, where re-modelling is required instead of re-tuning.

4.3 Oscillation-risk assessment from the measured characteristic

Finally, the measured dynamic characteristic is put to work in the characteristic-analysis stage of Fig. 2. The impedance-based stability criterion was applied with a Thevenin grid impedance Z_g = R_g + sL$_g$ (R_g = 0.1·ω$_0$L_g) swept over L_g ∈ {0.1, 0.3, 0.6, 1.2} mH, and the stability margin was defined as the minimum distance of the Nyquist curve $Z_g(f)/Z_l(f)$ to the critical point (−1, 0) on 2–470 Hz:

$$\eta = \left|1 + \frac{Z_g(f)}{Z_l(f)}\right| \tag{18}$$

Table *8* and Fig. 13 show that the margins computed from the 4VP-DFT impedance reproduce the analytic ranking with a worst-case deviation of 33% across the scan (31% at the critical Lg = 0.3 mH, where the margin dips to 0.045 and the system approaches the oscillation boundary), while margins from the plain-DFT impedance deviate by 84% at the critical point and misjudge the margin at Lg = 0.6 mH by a factor of six. This contrast is a direct manifestation of the error-amplification mechanism quantified in Section 2.1: because the sensitivity of the margin to the impedance grows as 1/η, the neighbourhood of the oscillation boundary—where risk assessment matters most—is exactly where impedance accuracy is most critical, and a measurement chain that is merely good on average is not sufficient there. Both measurements correctly flag the critical operating point, but only the 4VP-DFT provides the accuracy level needed for quantitative risk ranking—the exact capability required for online oscillation early warning of data-centre, storage and fast-charging interconnections [2,4,12].

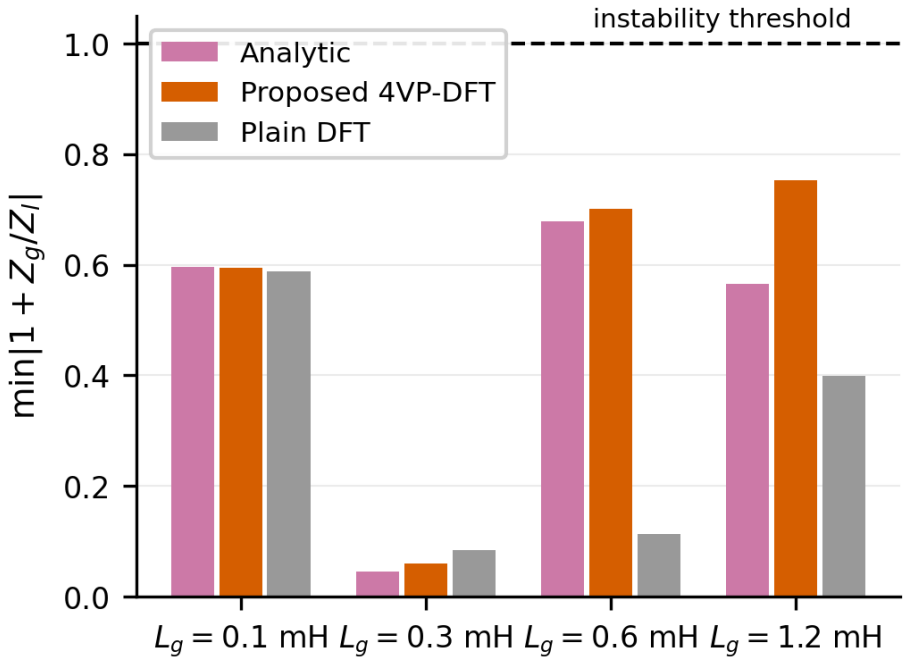


Figure 13. Impedance-based stability margin across the grid-inductance scan. The 4VP-DFT margins reproduce the analytic ranking and the critical point (L_g = 0.3 mH); plain-DFT margins deviate by more than 80% and misjudge the margin at L_g = 0.6 mH by a factor of six.

Table *8*. Impedance-based stability margin min|1 + Zg/Zl| for grid inductances Lg = 0.1–1.2 mH (Rg = 0.1ω_0L_g).

L_g (mH)	Analytic	From 4VP-DFT impedance	From plain-DFT impedance
0.1	0.596	0.594	0.588
0.3 (critical)	0.045	0.059	0.083
0.6	0.678	0.701	0.114
1.2	0.566	0.753	0.400

5. Discussion

The results support a clear division of labour between the two modelling routes. The model-driven route is structurally informative—it explains *why* the impedance peaks near 55 Hz (the DC-voltage loop interacting with the PLL)—but its parameters are nominal and drift in service; the data-driven route is parameter-free but is only as good as its phasor layer. Section 4.1 showed that the two routes agree within 1.0% when the measurement is accurate, and Section 4.2 showed that their discrepancy can be converted back into corrected model parameters by joint identification—recovering the PLL bandwidth within 2.8% and reducing the model-to-data discrepancy by a factor of six—provided the impedance fed into the identification is accurate. Because clean reference

waveforms yield exact parameter recovery, the residual identification error with 0.2 s windows is attributable to measurement noise rather than to weak identifiability. Measurement accuracy is thus not a detail of the data-driven route but the shared foundation of the entire modelling framework: it validates the model-driven route, constitutes the data-driven route, and decides whether the correction loop converges to the true parameters or to a wrong model.

Placed in the literature, the 4VP-DFT occupies a specific point in the estimator design space. The virtual-αβ method [35] and the offset-π/3 method [36] combine phasors from shifted windows within a 4N/3 observation span and cancel, respectively, one of the two leading leakage terms; the 4VP-DFT is the unique two-pair combination that cancels both the constant and the first-order term while keeping the equivalent noise bandwidth at 19/81 of the rectangular window. Windowed interpolated-DFT estimators [15,37] model the leakage explicitly and attain the highest static accuracy—machine precision on stationary tones in our benchmark—but require a two-cycle window, a frequency search and a higher sidelobe floor against on-bin harmonics, and they respond more slowly to steps. Taylor–Fourier estimators [14,38] generalize the model to time-varying phasors; the TF-2 baseline matches the 4VP-DFT within its three-cycle design band but degrades beyond ±5 Hz and is the most sensitive of the six estimators to wideband distortion (Table *4).* No existing estimator simultaneously offers single-cycle-class latency, second-order dynamic-error cancellation, an on-bin third-harmonic suppression of 1/9 and a noise bandwidth *6* dB below the rectangular window, which is the combination that the measurement layer of Fig. *2* requires.

For system operators the practical value is threefold. First, the static characteristic of Section 3.2 quantifies the negative-incremental-resistance strength of emerging loads under realistic workload fluctuations, an input to voltage-stability studies and protection settings near converter-dominated substations. Second, the wideband impedance of Section 4.1 enables impedance-based oscillation screening [8–11] without relying on confidential vendor models; the 33% worst-case deviation of the derived stability margin preserves the critical operating point and the ranking across grid strengths, which is sufficient for risk ranking. Third, and most directly tied to recent operational experience, the estimator directly addresses the measurement-adequacy gap identified for data-centre oscillation monitoring *[4]:* PMUs are effective at electromechanical frequencies but their phasor-estimation filters attenuate higher-frequency oscillation content. The 4VP-DFT replaces exactly that filtering stage inside existing PMU/WAMS chains, upgrading measurement points at data-centre, storage and fast-charging interconnections into external-characteristic monitors that feed oscillation early-warning systems [12]. The computational budget is compatible with existing PMU hardware: with the recursive DFT update the estimator costs five complex rotations per input sample—about 2×10^5 real multiply-accumulates per second at 9.6 kHz—plus five complex

additions per report, a small fraction of the capacity of current DSP and FPGA platforms that leaves the latency budget essentially unchanged. Accurate phasor-level *P* and *Q* estimates are equally the prerequisite for energy-flow-based oscillation-source location and damping evaluation [18], which needs trustworthy terminal power oscillations to attribute damping contributions. This capability is becoming urgent as wide-area oscillations induced by aggregated AI workloads are recognized as an operational risk *[2]* and data-centre consumption keeps growing *[3].*

Several limitations should be noted. *(i)* Under simultaneous amplitude–phase modulation all one-cycle-window estimators, including the 4VP-DFT, exhibit errors of several degrees because the upper modulation sideband coincides with the measurement frequency; multi-cycle or time-varying-frequency extensions are needed for strongly modulated workloads such as training-iteration power swings at high fluctuation rates [1,2]; a frequency-tracked 4VP-DFT, in which the virtual offsets are recomputed at the tracked frequency, is the most direct extension and would simultaneously remove the common-mode TVE floor of item (vi). (ii) The impedance scan characterizes the positive-sequence, small-signal behaviour only; negative-sequence and MIMO admittance characterization [11] requires additional injection patterns. (iii) The fundamental-removal stage assumes the grid frequency is known from a frequency-tracking stage; in strong electromechanical swings the removal and the impedance estimate should be computed in a sliding fashion, which is compatible with the estimator but was not exercised here. (iv) The validation was performed on a detailed simulation model cross-checked against its own Jacobian linearization (agreement within 1.0%); hardware-in-the-loop and field validation on a physical rectifier test bench—e.g. against an EMT reference model in the spirit of the open data-centre models [5]—is the natural next step; in the meantime, the Monte Carlo statistics of Table *4* define the accuracy envelope that such validation should reproduce, and a switching-level EMT cross-check of the plant model remains the main credibility gap of the present, purely model-based evidence. (v) The 4VP-DFT uses five N-point DFTs and a 5N/3-sample observation span per report (33.3 ms at 50 Hz versus 20 ms for the plain DFT); for scanning applications, where windows are processed off-line or at low duty cycle, this overhead is acceptable, and with recursive DFT updating the incremental cost is five complex rotations per input sample—negligible next to the plant-model integration (Methods). (vi) The residual TVE under large frequency offsets is a common-mode floor (2.27–2.37% at −5 Hz across all compared estimators, Table *2)* set by the rectangular window's magnitude response and the reference-phasor convention, rather than by the dynamic leakage error targeted in this paper; amplitude-domain compensation would be a complementary extension. (vii) Interferers closer than about 25 Hz to the fundamental lie inside the main lobe of every compared window: neither the 4VP-DFT nor the Hann-ipDFT baselines reject such near-band tones, so the known-frequency fundamental-removal stage of Section 4.1—or a band-limiting pre-filter—remains essential for

impedance scanning; a second, tone-matched virtual-phasor bank is a possible estimator-level extension.

Taken together, these points show that the estimator, the measurement chain and the modelling loop cannot be designed independently: the value of every downstream analysis is capped by the accuracy of the phasor layer that feeds it.

6. Conclusion

This paper addressed the measurement-distortion bottleneck that limits data-driven modelling of the external characteristics of new-type power-electronic loads such as data-centre rectifier fronts, storage PCS and fast chargers. A four-virtual-phasor DFT estimator was derived that combines five windows with phase-shifted replicas to cancel the mirror-component leakage of single-bin DFT phasor estimation, and its design rules and leakage-rejection properties were established analytically and summarized as three propositions. Compliance tests extended to ±10 Hz offsets, 10% THD, interharmonic interference and 200-realization Monte Carlo statistics, and benchmarked a*g*ainst interpolated-DFT and second-order Taylor–Fourier baselines, showed that the estimator surpasses every single-cycle baseline in phase accuracy under frequency offset, harmonic distortion, noise and out-of-band interference—reducing the best baseline's maximum phase error by 37% at a −5 Hz offset and by up to *41%* under interference—while its 40 ms response remains within the 70 ms grid-code bound; the residual TVE under off-nominal frequency is a common-mode floor shared by all rectangular-window estimators (Discussion).

On a *2* MVA rectifier-type load with drifted control parameters, the estimator made both faces of the external characteristic measurable: the static V–I trajectory exposed the negative incremental resistance of the constant-power behaviour, and the wideband impedance measured in 0.2 s windows coincided with the 10 s ground truth, while the ground truth and the Jacobian linearization of the average model—the data-driven and model-driven references—agreed with each other within 1.0% at every scan tone. The same measured impedance closed the loop between the two routes: joint identification recovered the drifted PLL and current-loop bandwidths to within 2.8% and 10.9% and reduced the model-to-data deviation from 55.0% to 8.73% (a factor of 6.3), whereas the same procedure driven by plain-DFT measurements converged to a wrong model—demonstrating that measurement accuracy, not identifiability, decides whether model correction succeeds. A 36-scenario robustness scan over the nominal-parameter plane showed that the correction loop retains a median post-correction deviation of 11.0% across single and joint drifts, degrades gracefully to a clearly detectable saturation when the drift exceeds the retuning authority, and should be gated by the measurement uncertainty to avoid corrupting an already accurate model. Finally, the

impedance-based stability margins derived from the measured characteristic reproduced the analytic oscillation-risk ranking across grid strengths with a worst-case deviation of 33%, preserving the critical operating point that plain-DFT-based margins misjudge by up to a factor of six.

By treating the measurement algorithm—not the injection hardware or the sensor—as the design variable, the proposed method turns the external-characteristic measurement of converter-dominated loads into a standard-compliant, online measurement function and lets field measurements keep physics-based load models accurate as equipment and firmware evolve. Hardware-in-the-loop and field validation on a physical rectifier test bench, extension to negative-sequence and MIMO admittance characterisation, and tracking of strongly modulated workloads are the natural next steps.

7. Methods

Load model: average model of the grid-feeding power-electronic load

The load under test is a three-phase grid-feeding voltage-source converter with an *L* filter supplying a constant-power DC load—an average representation of data-centre rectifier fronts, storage PCS and EV fast chargers [5,6]. The model is formulated in the dq frame with state vector [id, iq, vdc, xiv, xid, xq, θpll, xpll, vcd, vcq] and includes a synchronous-reference-frame PLL (bandwidth 10 Hz in the as-built configuration; 15 Hz as the nominal design value in the correction experiment of Section 4.2), inner current loops (bandwidth 400 Hz in the as-built configuration; 300 Hz as the nominal design value, with decoupling and a 1/6000 s modulation lag), a DC-link voltage loop (kp = 40, ki = 400) and the constant-power DC load P(t) on a 20 mF DC-link capacitor. Parameters are listed in Table *9*. The PCC voltage is stiff; perturbations are injected as a balanced positive-sequence set. A Jacobian linearization of the same model at the operating point provides the analytic impedance reference $Z_{an}(f) = 2/[G_{00} + G_{11} + j(G_{10} - G_{01})]$ with $G(s) = C(sI - J)^{-1}B$, where the input vector is the injected dq voltage pair and the output the grid-frame converter current; the same linearization, parameterized in both the PLL and current-loop bandwidths *(β = (fpll, fcur),* with the proportional–integral gains mapped as $K_p = \sqrt{2}\omega_n L_f$, $K_i = \omega_n R_f$ for a bandwidth $\omega_n/2\pi$), serves as $Z_{model}(f; \beta)$ in Eq. *(17*). The joint least-squares search of Eq. (16) first evaluates the residual on a coarse grid (PLL 2–30 Hz in 0.5 Hz steps, current loop 100–800 Hz in 20 Hz steps) and then refines the best point by coordinate descent with halving steps down to 0.02 Hz / 0.8 Hz.

Table *9*. Main parameters of the grid-feeding converter load (average dq model).

Parameter	Value	Parameter	Value

Rated power S_b	2 MVA	DC-link capacitance C_dc	20 mF
CPL power P_0	1.6 MW	DC voltage setpoint V_dc	1200 V
PCC voltage (L-L)	690 V	PLL bandwidth (as built)	10 Hz
Filter L_f / R_f	0.15 mH / 0.01 Ω	Current-loop bandwidth	400 Hz (as built)
Modulation lag T_d	1/6000 s	DC-voltage loop k_p / k_i	40 / 400

Impedance-scan measurement procedure

A balanced positive-sequence multi-sine voltage perturbation with 16 tones (25–475 Hz, *5* Hz spacing on the 0.2 s window grid, 2% per-tone amplitude, randomized phases) is superimposed at the PCC. Terminal waveforms are sampled at 9.6 kHz. For each 0.2 s analysis window the 49.5 Hz fundamental is removed by the least-squares fit of Eq. *(15*) over a span of 1.2 s surrounding the window; the residual is then analysed at each tone bin with the four estimators of Table *1,* and the impedance is formed as Eq. *(16*) and averaged over ten windows. The ground truth is the Hann-windowed DFT of the 10 s clean record at the tone bins (rectangular-window leakage of the 563 *V* fundamental would otherwise contaminate the truth itself at the low-frequency tones); the analysis copy is stressed with 2%/1.2% fifth/seventh harmonics at half-bin offsets and 50 dB white noise.

Compliance tests and simulation platform

The compliance tests of Section 3.1 follow the steady-state and dynamic test-signal definitions of the Chinese synchrophasor testing standard [17] with N = 192: frequency offsets ±1/±3/±5 Hz, extended to ±10 Hz in Section 3.1; a 5.5% THD harmonic mixture (2nd–25th) with 50 dB noise; out-of-band tones at f_0 ± 25 Hz (the standard swept single-tone test), complemented by dual tones at f_0 + 75 Hz and f_0 + 125 Hz and an interharmonic at f_0 + 125.5 Hz at 10% amplitude; simultaneous amplitude–phase modulation (ka = kp = 0.1, fm = *5* Hz); and a 10% magnitude / 10° phase step for the response-time test, whose 70 ms limit is specified in [16]. All noisy conditions were repeated over 200 Monte Carlo realizations with randomized harmonic phases, and errors are reported as the maximum over an eight-point initial-phase grid per realization with the mean, standard deviation and 95% confidence interval across realizations (Table *4).* All simulations were implemented in Python: the load model was integrated with a fixed-step fourth-order Runge–Kutta scheme at 38.4 kHz and decimated to 9.6 kHz, and the estimators operate on the decimated waveforms. The 10 s impedance-scan run takes about one minute on a desktop computer, dominated by the RK4 integration rather than by the estimators.

Declarations

Data availability. The datasets generated and analysed during the current study are available from the corresponding author upon reasonable request.

Code availability. The Python implementation of the estimators, the load model and the test scripts are available from the corresponding author upon reasonable request.

Funding. This work was supported by the National Natural Science Foundation of China (NSFC) under Grant 52637007.

Author contributions. L.C. conceived the study, developed the 4VP-DFT estimator, performed the simulations, and wrote the original draft. M.Y.Z. contributed to the theoretical analysis, simulation design, and algorithm validation. F.X. contributed to the impedance measurement methodology, model-correction framework, and oscillation-risk assessment. C.X. contributed to the load modelling, stability analysis, and results interpretation. M.Z. supervised the work, secured resources, and validated the engineering applicability. All authors reviewed and edited the manuscript.

References

[1] Choukse, E., Warrier, B., Heath, S. et al. Power stabilization for AI training datacenters. arXiv:2508.14318 (2025).

[2] Ko, M.-S. & Zhu, H. Wide-area power system oscillations from large-scale AI workloads. arXiv:2508.16457 (2025).

[3] Shehabi, A. et al. 2024 United States Data Center Energy Usage Report. Report LBNL-2001637, Lawrence Berkeley National Laboratory (2024).

[4] Chatterjee, K., Follum, J., Varghese, A., Biswas, S., Farantatos, E. & Zhu, L. Measurement adequacy for monitoring data center oscillations. Report PNNL-39191, Pacific Northwest National Laboratory (2026).

[5] Wang, X. et al. Dynamic modeling of AI data center load in PSCAD. Version 1.0 report, Electric Reliability Council of Texas (ERCOT), 14 May 2026.

[6] Chen, Y., Zhao, Y., Geng, G. et al. Suppression strategy of ultra-low frequency oscillation in Yunnan power grid with BESS. Journal of Modern Power Systems and Clean Energy 9, 656–663 (2021).

[7] Yang, N., Ma, W., Wang, X., Xie, D. & Gu, C. Defining SSO power and characterizing SSO propagation in power systems with wind farms integration. IEEE Transactions on Power Systems 36, 3531–3540 (2021).

[8] Middlebrook, R. D. Input filter considerations in design and application of switching regulators. In Proc. IEEE Industry Applications Society Annual Meeting, 366–382 (1976).

[9] Sun, J. Impedance-based stability criterion for grid-connected inverters. IEEE Transactions on Power Systems 26, 307–315 (2011).

[10] Céspedes, M. & Sun, J. Impedance modeling and analysis of grid-connected voltage-source converters. IEEE Transactions on Power Electronics 29, 1254–1261 (2014).

[11] Rygg, A., Molinas, M., Zhang, C. & Cai, X. A modified sequence-domain impedance definition and its equivalence to the dq-domain impedance definition for the stability analysis of AC power electronic systems. IEEE Journal of Emerging and Selected Topics in Power Electronics 4, 1383–1396 (2016).

[12] Phadke, A. G. & Thorp, J. S. Synchronized Phasor Measurements and Their Applications (Springer, 2008).

[13] Benmouyal, G. An adaptive sampling-interval generator for digital relaying. IEEE Transactions on Power Delivery 16, 624–633 (2001).

[14] de la Serna, J. Á. Dynamic phasor estimates for power system oscillations. IEEE Transactions on Instrumentation and Measurement 56, 1648–1657 (2007).

[15] Grandke, T. Interpolation algorithms for discrete Fourier transforms of weighted signals. IEEE Transactions on Instrumentation and Measurement IM-32, 350–355 (1983).

[16] IEEE Std C37.118.1-2011: IEEE Standard for Synchrophasor Measurements for Power Systems—Part I: Measurement (IEEE, 2011).

[17] GB/T 26862-2011: Testing specification of synchrophasor measurement units for power systems (Standardization Administration of China, 2011). (in Chinese)

[18] Chen, L., Min, Y., Chen, Y. et al. Evaluation of generator damping using oscillation energy dissipation and the connection with modal analysis. IEEE Transactions on Power Systems 29, 1393–1402 (2014).

[19] Y.Gao, Z. Dou, C. Zhang, C. Jiang and L. Wang, "Research on adjustable load classification and aggregation and external characteristic identification method," 2023 8th International Conference on Power and Renewable Energy (ICPRE), Shanghai, China,1449-1454 (2023).

[20] Z.Wu et al., "Analysis of External Short-Circuit Fault Characteristics and Safety Impact in PEMFC System Integrated With Renewable Energy," in IEEE Transactions on Sustainable Energy, vol. 17, no. 2, 1715-1727 (2026).

[21] Q. Wu, S. Shen, H. Zhang and Q. Wang, "High Accuracy and Wide Temperature Range Converter-level On-State Voltage Measurement With Localization Function for SiC MOSFETs," in IEEE Transactions on Power Electronics, vol. 40, no. 4, 5971-5982 (2025).

[22] Z. Liu, J. Liu and Z. Liu, "A Sampling Optimized Method for Frequency Response Measurement in Power Electronics Systems," in IEEE Transactions on Power Electronics, vol. 37, no. 7,8447-8464 (2022).

[23] X. Xu, K. Peng, X. Liu, H. Wang and X. Fan, "High-Accuracy Cylindrical Capacitive Sensors Designed According to a Posture Error Model for the Combined Measurement of Linear and Angular Displacements," in IEEE Transactions on Instrumentation and Measurement, vol. 75, 1-13 (2026).

[24] Z. Su and J. Du, "Optimization of External Characteristic Parameters of Multiple Motors for Electric Vehicles Based on Optimal Torque Distribution Coefficient," 2022 25th International Conference on Electrical Machines and Systems (ICEMS), Chiang Mai, Thailand, 1-5 (2022).

[25] Huijie-Lin, X. Yang, B. Yin and X. Wang, "Coupled Optimization of Internal and External Battery Characteristics for Grid-Connected Electric Vehicles," 2026 5th International Conference on Smart Grid and Green Energy (ICSGGE), Dongfang, China, 706-712 (2026).

[26] Guo et al., "Research on Flexible Load Control Strategy and Power Supply Reliability Management Based on Sensitivity Analysis in New Type Power System," 2024 7th International Conference on Renewable Energy and Power Engineering (REPE), Beijing, China, 220-224 (2024).

[27] Portella, F. M. Bayer and R. J. Cintra, "Multiplierless DFT Approximation Based on the Prime Factor Algorithm," in IEEE Transactions on Signal Processing, vol. 73, 5273-5285 (2025).

[28] Wang, H. Wen, G. Li and L. Wang, "Iterative Two-Point Interpolated DFT Algorithm for Accurate Frequency Estimation," in IEEE Transactions on Power Delivery, vol. 37, no. 3, 1671-1681 (2022).

[29] Wang, Y. Liang, Y. Zhang and Z. Zhu, "Multi-Term Cosine-Sum Windows-Assisted DFT-IDFT-Based Minimum-Segment Calibration for SAR ADCs," in IEEE Transactions on Circuits and Systems , vol. 73, no. 3, 1683-1696 (2026).

[30] Liu, L. Fan, H. Wu, J. Jin, J. Xing and Z. Liu, "DFT-Based Frequency Estimation of Multiple Sinusoids," in IEEE Access, vol. 10, 40230-40236 (2022).

[31] Tang, J. He, K. Wang, B. Wei and X. Tang, "DFT-Based Method for Accurate I/Q Imbalance Estimation," in IEEE Transactions on Circuits and Systems vol. 71, no. 4, 2474-2478 (2024).

[32] S. Dorji, G. Nourbakhsh, Y. Mishra and Y. Zhang, "Impact of Increasing Level of Power Electronic Load and Distributed PV on Power System Stability," 2025 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe), Valletta, Malta, 1-5 (2025).

[33] G. Ma, B. Luo, J. Luo, G. Zhu and J. Lu, "An energy-fed three-phase four-wire power electronic load compatible with multiple load requirements," 2023 10th International Forum on Electrical Engineering and Automation (IFEEA), Nanjing, China, 2023, 282-287 (2023).

[34] Y. Tao, J. Jiang, X. Lu, W. Zhang and C. Zhang, "Improved Series Arc Model for Power Electronic Loads," in IEEE Transactions on Plasma Science, vol. 53, no. 11, 3476-3485 (2025).

[35] Zhou, N. & Liao, Y. A phase angle measurement algorithm based on a virtual αβ stationary coordinate. Automation of Electric Power Systems 36, 61–65, 76 (2012). (in Chinese)

[36] Zhang, L., Xu, K., Xu, H. et al. Partial π/3 algorithm of synchronized phasor angle measurement for anti-harmonic interference. Automation of Electric Power Systems 43, 149–155 (2019). (in Chinese)

[37] Harris, F. J. On the use of windows for harmonic analysis with the discrete Fourier transform. Proceedings of the IEEE 66, 51–83 (1978).

[38] Platas-Garza, M. A. & de la Serna, J. A. Dynamic harmonic analysis through Taylor–Fourier transform. IEEE Transactions on Instrumentation and Measurement 60, 804–813 (2011).