

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/189829/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Khalid, Junaid, Uddin, Muhammad Helal, Smailes, Michael, Liang, Jun and Wang, Sheng 2026. Physics informed adaptive frequency support with rotor energy-aware recovery for grid forming wind turbines. IEEE Transactions on Sustainable Energy 10.1109/tste.2026.3737065

Publishers page: <https://doi.org/10.1109/tste.2026.3737065>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Physics Informed Adaptive Frequency Support with Rotor Energy-Aware Recovery for Grid Forming Wind Turbines

Junaid Khalid, *Student Member, IEEE*, Muhammad Helal Uddin, *Student Member, IEEE*, Michael Smailes, *Member, IEEE*, Jun Liang, *Fellow, IEEE*, Sheng Wang, *Senior Member, IEEE*

Abstract—High penetration of converter-interfaced wind power reduces system inertia and primary frequency response. This makes frequency nadir and rate-of-change-of-frequency (RoCoF) increasingly difficult to control. While kinetic-energy (KE)-based synthetic inertia from wind turbines can provide fast support without long-term energy loss. However, existing designs rely on fixed or heuristic gain scheduling and ad hoc rotor-speed recovery, which may underutilise available KE, cause unsafe rotor deceleration, or lead to secondary frequency dips. This paper proposes a physics-informed, constraint-embedded neural adaptive (PINA) framework for grid-forming wind turbines. The machine-side converter provides KE-based inertia and droop response, while the grid-side converter operates in virtual synchronous control mode to regulate the dc-link voltage and maintain grid synchronisation. An adaptive gain network (Gain-Net) schedules inertia and droop gains using local measurements of frequency deviation, RoCoF, rotor speed, and normalised rotor-KE headroom, while RecNet shapes energy-aware rotor-speed recovery. Both networks are trained with physics-informed objectives and deployed with hard rotor-energy, converter-power, and ramp-rate constraints. Numerical case studies evaluate the proposed framework against recent state-of-the-art frequency-support methods.

Index Terms—Physics-informed constraint learning, frequency support, grid-forming wind turbines, rotor kinetic energy.

I. INTRODUCTION

WIND turbines (WTs) are typically operated in maximum power point tracking (MPPT) mode to maximise wind energy extraction. While effective for energy efficiency, MPPT decouples turbine active power from system frequency and limits their contribution to frequency regulation [1]. As converter-interfaced wind generation increasingly replaces synchronous generators (SGs), the aggregate synchronous inertia and governor response decline, making frequency excursions harder to arrest. The frequency nadir (FN) is a key indicator of frequency stability; if it falls below protection thresholds, under-frequency load shedding can be triggered, risking large-scale outages [2]. The 2019 Great Britain blackout, where insufficient frequency response from renewables was one contributing factor [3], highlights the

need to enable WTs to provide inertia support and primary frequency response.

To address this, three main sources of additional active power have been explored: (i) de-loaded wind operation [4], (ii) integrated energy storage systems (ESS) [5], and (iii) the kinetic energy (KE) stored in turbine rotors. De-loading creates active-power headroom by operating below the MPPT curve, typically via pitch-angle adjustment [6] or rotor overspeed control [7], but sacrifices energy yield. ESS can deliver fast and flexible support, yet adds cost and complexity. KE-based approaches maintain MPPT during normal operation and temporarily release or absorb rotor KE during disturbances, providing fast frequency support without long-term energy or economic loss. Therefore, this work focuses on KE-based frequency support.

The feasibility of KE-based frequency support from variable-speed WTs was established early by emulating SG inertial response through a RoCoF-proportional power term [8] and by characterising the inherent frequency behaviour of doubly-fed and fixed-speed machines [9]. Because WTs in MPPT mode do not inherently regulate active power when grid frequency deviates, additional control loops are required to extract KE and provide frequency support [10]. A common approach is to emulate SG inertia and droop behaviour by feeding back the rate of change of frequency (RoCoF) and frequency deviation into the active-power reference [11]. The resulting response naturally splits into a support phase, where KE is released to provide extra power, and a recovery phase, where rotor speed is restored while avoiding a secondary frequency dip (SFD).

However, designing effective KE-based support is challenging for several reasons. First, WT operating states within a wind farm are heterogeneous, so their releasable KE and safe deceleration margins differ significantly. Fixed-gain synthetic inertia can therefore under-utilise high-speed turbines and over-decelerate low-speed ones. To improve utilisation, time-varying or state-dependent gains have been proposed [12]. For example, [13] schedules inertia and droop gains with rotor speed, while [14] initially boosts gains and then adjusts them based on RoCoF and rotor speed as the system approaches the frequency nadir. However, these schemes still rely on heuristic tuning and only partially exploit the joint dependence on grid dynamics and turbine state.

Second, the duration of frequency support is as important as its initial magnitude. Before the grid reaches FN, WTs should continually provide frequency support by delivering extra power beyond their initial MPPT points [15]. The time to nadir is driven by system inertia and the aggregate extra energy

J. Khalid, M.H. Uddin, and J. Liang (corresponding author, email: LiangJ1@cardiff.ac.uk) are with the School of Engineering, Cardiff University, CF24 3AA Cardiff, U.K.

M. Smailes is with Offshore Renewable Energy Catapult, NE24 1LZ Blyth, Northumberland, U.K.

S.Wang (corresponding author, email: Sheng.Wang@glasgow.ac.uk) is with the James Watt School of Engineering, University of Glasgow, G12 8QQ Glasgow, U.K.

This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/X031144/1]. Project title: Accelerating the deployment of offshore wind using DC technology (ADoreD).

contributed by WTs and other resources, whereas the support duration of each WT depends on its physical characteristics, operating point, and control strategy [16]. As KE release slows the rotor, turbines starting at lower speeds may reach their minimum-speed limits before the nadir, triggering rotor-speed protection and prematurely removing support, which can worsen the FN [17]. Various rotor-recovery schemes have been proposed, including washout filters in the droop loop [18], stepwise inertial control with staged recovery [19], and fuzzy logic to avoid SFDs [20]. Washout filters can reduce sustained support; stepwise schemes may still cause SFD due to abrupt power reductions; and fuzzy controllers rely on rule bases that may not generalise outside nominal conditions. Therefore, existing designs rely on fixed or heuristic gain scheduling and heuristic rotor-speed recovery, which may not fully exploit available KE and may cause unsafe rotor deceleration or SFD.

Beyond fixed, scheduled, and heuristic gain designs, several advanced control paradigms have been investigated. At the converter level, the inertial response can be accelerated by embedding inertia and damping directly in the synchronisation loop, as in the inertial phase-locked loop (IPLL) [21], or by feeding frequency deviation forward into the PLL (FFDC) [22]. These methods deliver fast response, but they are largely formulated for DFIG or grid-following (GFL) structures in which PLL dynamics are central. They primarily reshape how the frequency signal is sensed or injected, rather than how limited rotor KE is allocated and recovered. Adaptive temporary frequency support (ATFS) schemes for DFIGs address recovery explicitly by setting the initial power reference from the initial RoCoF and shaping the support and recovery curves [23]. However, they still depend on event-specific and predesigned power-trajectory rules. Model predictive control (MPC) can optimise support over a prediction horizon while explicitly handling constraints [24], but requires repeated online optimisation and an accurate predictive model, which is computationally demanding for fast turbine-level implementation.

Learning-based methods have also been applied to frequency support. Reinforcement learning (RL) can discover nonlinear control policies directly from data, including adaptive inertia and droop gains for wind farms [25]. However, the resulting policies are typically black-box, and safe exploration, constraint satisfaction, and interpretability remain difficult during deployment. Physics-informed learning incorporates physical knowledge through governing-equation residuals, physical constraints, or physically structured inputs and outputs [26]. In power and energy systems, physics-informed neural networks (PINNs) have been used mainly for analysis and identification, such as dynamic simulation and parameter estimation [27], optimal power flow with feasibility regularisation [28], and physically consistent wind-farm aggregation [29]. In these works, the network output is usually a predicted trajectory, an estimated parameter, or a system model. In this paper, by contrast, physics-informed learning is used in a controller-design sense. Physical knowledge enters through structured input features, bounded inertia and droop outputs, hard support-power and rotor-energy constraints, slew-rate limits, and swing-equation and power-consistency residuals.

Thus, the neural networks provide nonlinear adaptation, while the deployed controller remains interpretable and feasible by construction. This distinction is important because residual penalties alone do not guarantee that the deployed controller will respect rotor-speed, converter-power, and ramp-rate limits.

A further layer of complexity arises from the dc-link and grid-forming (GFM) behaviour of the WT converter. Compared with GFL control, GFM control lets the WT behave as a controlled voltage source and improves dynamic performance in low-inertia systems, but it requires active-power freedom at the dc terminals supported by stable dc-link voltage [30]. Several studies regulate the dc-link from the machine-side converter (MSC) in GFM-WT configurations [31], but this may adversely affect the mechanical parts of WTGs [32]. To address this, [33] proposed a virtual synchronous control (ViSynC) strategy in which the grid-side converter (GSC) uses dc-link voltage squared as a synchronising state, effectively leveraging dc-link energy to support grid synchronisation and emulate a damping-like response. Based on this control strategy, [34] integrated a supercapacitor energy storage system (SC-ESS) on the dc-link to supply additional power during frequency events. More recently, [35] modelled rotor KE and SC-ESS as combined equivalent capacitors in parallel with the dc-link to coordinate frequency support. These GSC-centric schemes improve dc-link stability, but rely on predefined capacitance values, energy-sharing ratios, and control gains. Moreover, they largely treat the MSC as a passive energy source, underutilising its direct access to rotor dynamics for KE management.

In summary, although KE-based frequency support has advanced considerably, no existing approach simultaneously: (i) adapts inertia and droop action to both event severity and rotor-energy headroom; (ii) enforces rotor-energy, converter-power, ramp-rate, and gain-envelope feasibility in the deployed controller; (iii) coordinates post-nadir recovery without SFD; and (iv) does so within a GFM Type-IV architecture that separates rotor-KE support from dc-link and grid-forming stabilisation. The remaining challenge is therefore not simply to make inertia or droop gains adaptive, but to make the adaptation state-aware, physically feasible, interpretable, and deployable in real time. ATFS and heuristic AVIC methods improve recovery behaviour, but their support and recovery trajectories remain largely predesigned. MPC can handle constraints explicitly, but requires online optimisation and an accurate prediction model. RL can learn nonlinear policies, but the resulting action is usually difficult to interpret and does not inherently satisfy rotor-speed, converter-power, and ramp-rate limits. Classical PINN-based methods mainly use physics residuals to learn trajectories, parameters, or surrogate models rather than to construct a hard-constrained frequency-support controller.

To address these gaps, this paper proposes PINA, a physics-informed, constraint-embedded neural adaptive framework for Type-IV GFM wind turbines. PINA uses physics-informed learning in a control-oriented manner. The neural networks are used to schedule interpretable inertia/droop and recovery parameters, rather than to generate an unconstrained power command. These learned parameters are deployed inside a

physically constrained control structure, so the final command remains bounded by rotor-energy headroom, converter-power limits, ramp-rate limits, and hard inertia-droop gain envelopes. The framework is implemented with additional MSC and GSC roles, allowing rotor-KE-based frequency support and dc-link/grid-forming stabilisation to be coordinated without inertia double-counting. The main contributions of this article are as follows:

- A GFM Type-IV WT frequency-support architecture is presented with explicit converter-role separation. The MSC provides rotor-KE-based frequency support via adaptive inertia/droop action, while the GSC operates in ViSynC mode for dc-link regulation and synchronisation, with no explicit virtual inertia. This reduces loop interactions and avoids inertia double-counting.
- A physics-informed, constraint-embedded GainNet is proposed for adaptive inertia/droop scheduling. Rather than using a fixed gain, a one-dimensional rotor-speed schedule, or a heuristic rule base, GainNet maps the joint event-turbine state $(\Delta f, \dot{f}, \omega_r, \mu)$ to bounded inertia and droop gains. This allows aggressive support when event severity and rotor-energy headroom are both high, and conservative support near rotor-speed or converter limits.
- A RecNet-based recovery mechanism is introduced to restore rotor speed after the nadir without creating a secondary frequency dip. RecNet outputs bounded recovery parameters, while the final recovery command is shaped by post-nadir gating, energy-to-target capping, tapering near the pre-event rotor speed, and ramp-rate limits.
- A deployment-safe learning pipeline is established in which offline physics-informed training is combined with online deterministic inference and hard feasibility projection. Unlike black-box learned policies or on-line optimisers, the deployed controller remains an interpretable inertia/droop-plus-recovery law constrained by rotor-energy, converter-power, ramp-rate, and hard gain-envelope limits. Reproducible code is available on [GitHub](#).

II. FREQUENCY SUPPORT MECHANISM AND RATIONALE BEHIND THE PROPOSED DESIGN

A. Kinetic-Energy-Based Frequency Support

In normal operation, a variable-speed WT operates in MPPT mode and its active power does not explicitly track grid frequency. During a frequency event, however, the turbine can temporarily deviate from MPPT and inject an additional active-power component that draws on rotor kinetic energy.

Let f denote the grid frequency, f_{nom} the nominal frequency, and define the frequency deviation and its rate of change (RoCoF) as

$$\Delta f(t) \triangleq \frac{f(t) - f_{\text{nom}}}{f_{\text{nom}}}, \quad \dot{f}(t) \triangleq \frac{d}{dt} \Delta f(t) \quad (1)$$

Let ω_r be the mechanical rotor speed, P_m the mechanical power extracted from the wind, and P_e the electrical power delivered to the grid. The electro-mechanical balance of the drivetrain is




Fig. 1. Comparison of conventional designs: Fixed IC and PIC over-decelerate the rotor, and abrupt power recovery triggers an SFD.

$$2H_w \omega_r \dot{\omega}_r = P_m - P_e \quad (2)$$

where H_w is the equivalent inertial time constant of the WT. A widely used synthetic-inertia and droop (PD-type) support law adds an auxiliary power signal

$$\Delta P(t) = -K_d \dot{f}(t) - K_p \Delta f(t) \quad (3)$$

to the MPPT power reference. Here K_d and K_p are the inertial and droop gains, respectively. The resulting active-power reference is

$$P_{\text{ref}}(t) = P_{\text{MPPT}}(\omega_r(t)) + \Delta P(t) \quad (4)$$

where $P_{\text{MPPT}}(\omega_r) = k_{\text{opt}} \omega_r^3$ represents the MPPT curve at the current rotor speed. During a support interval $[t_0, t_1]$, the net rotor kinetic energy released to the grid is

$$\Delta E_w = \int_{t_0}^{t_1} (P_e(t) - P_m(t)) dt = H_w (\omega_r^2(t_0) - \omega_r^2(t_1)) \quad (5)$$

so the maximum sustainable support is fundamentally limited by rotor speed headroom between the initial speed and the minimum allowable speed. In practice, the auxiliary power must also satisfy converter and mechanical limits, e.g.

$$P_{\text{min}} \leq P_{\text{ref}}(t) \leq P_{\text{max}}, \quad \left| \dot{P}_{\text{ref}}(t) \right| \leq \gamma \quad (6)$$

where P_{max} is often slightly above rated power and γ bounds the ramp rate. Similar envelopes apply to rotor speed and torque.

Equations (2)–(5) show that the effective energy reservoir and torque stress during support depend strongly on the instantaneous rotor speed. At the same time, the impact of (K_d, K_p) on system-wide frequency is different: K_d primarily shapes RoCoF by emulating inertia, whereas K_p mainly improves the frequency nadir through a droop-like response, with limited effect on the final steady-state frequency because the rotor energy must ultimately be recovered.




Fig. 2. Proposed framework with ViSynC on the GSC and KE-based frequency support on MSC using adaptive inertia and droop gains provided by PINA.

B. Limitations of Conventional Designs

Existing synthetic-inertia and droop designs based on rotor KE can be broadly grouped as follows:

- Fixed-gain PD: constant (K_d, K_p) are tuned to provide acceptable RoCoF and FN under a representative disturbance. This approach ignores the heterogeneity of turbine operating conditions: high-speed machines under-utilise their KE reserve, while low-speed machines may be over-decelerated during large events.
- Rotor speed-scheduled gains: to reflect available KE, several works schedule K_d and K_p as functions of ω_r , often with quadratic or piecewise-linear profiles [12]–[14]. Gains increase when more KE is available and decrease as the rotor slows down. However, these schedules are typically derived offline and only partially account for disturbance severity as measured by Δf and $\dot{f}$.
- Fuzzy and heuristic adaptive schemes: fuzzy logic or rule-based controllers adjust droop and inertia gains using combinations of $|\Delta f|$, $\dot{f}$, and rotor-related signals [15], [20]. Their performance depends heavily on the rule base and membership functions, which are difficult to tune and may not generalise across grids and operating points.

After the nadir, rotor speed must be recovered without

provoking an SFD. Common strategies include: (i) PI-based continuous recovery that subtracts a negative power offset proportional to rotor speed error; (ii) washout-filter-based droop, where low-frequency components of Δf are gradually removed so that the WT naturally returns to MPPT [18]; and (iii) stepwise recovery, where P_{ref} is reduced by a preset step after a fixed delay [19]. All involve a trade-off between fast rotor recovery and avoiding SFD and are sensitive to gain and timer settings. A typical nadir detector enables the recovery module once $\Delta f(t) \dot{f}(t) < 0$ persists for a short dwell time, indicating that the frequency has passed its minimum. Fig. 1 illustrates the limitations of conventional designs and compares MPPT operation with fixed-gain (Fixed IC), adaptive rotor-speed-scheduled (AIC), and piecewise-linear-scheduled (PIC) inertia control under an identical load disturbance. These results confirm that conventional designs face an inherent trade-off between aggressive support (over-deceleration and SFD) and conservative support (under-utilised KE).

In summary, conventional designs do not simultaneously capture: (i) the nonlinear dependence of optimal support on rotor speed headroom, (ii) the severity and time evolution of the frequency event, and (iii) converter and mechanical constraints. Manually designed gain schedules and recovery laws become increasingly difficult to tune as wind penetration and

TABLE I
POSITIONING OF PINA RELATIVE TO EXISTING FREQUENCY-SUPPORT CONTROL METHODS

Method	Adaptation mechanism	Constraint handling	Deployment feasibility	Online burden	Technical gaps addressed by PINA
ATFS [23]	Predefined event dependent power trajectory	Trajectory limits; event-set ceiling	Limited by trajectory envelope	Low	Predefined support shape
MPC [2], [24]	Online optimisation	Explicit in optimisation	Only within the model	High; repeated optimisation	Computational burden; model dependence
RL / DRL [25]	Reward-based policy	Not inherent; training-dependent	Not guaranteed by action structure	Low inference; high training	Black-box action; interpretability concern
Classical PINN [26]–[29]	Residual-trained network	Not deployment-enforced	Residual alone doesn't guarantee feasibility	Low inference	No deployed control action
Proposed PINA	Offline-learned GainNet/RecNet scheduling maps	Hard power, energy, slew rails	High command-level feasibility	Deterministic inference; 2.9 μs	Adaptive, interpretable, feasible control

grid conditions evolve. Furthermore, the limitations of more advanced model-based (e.g. MPC) and learning-based (e.g. RL) methods were discussed in Section I, and a quantitative comparison of the proposed controller against representative recent techniques is presented in Section V.

C. ViSynC-Based Grid-Forming Control

Grid-forming control addresses the need for stable operation in weak grids and during large disturbances. The ViSynC scheme regulates the square of the dc-link voltage as a synchronising state for the GSC. Its frequency reference can be written as

$$\frac{\omega}{\omega_0} = \frac{\omega_g}{\omega_0} + \frac{s + K_t}{K_j s + K_d} \frac{V_{DC}^2 - V_{DC}^{*2}}{V_{DC0}^2} \quad (7)$$

where ω_g is the grid frequency, K_t is the dc-link tracking gain, K_j and K_d are control coefficients, V_{DC} is the dc-link voltage with reference V_{DC}^* , and V_{DC0} is a nominal value. Together with the dc-link power balance

$$\frac{C}{2} s V_{DC}^2 = P_{DC} - P_{AC} \quad (8)$$

and the AC power angle relation, (7) yields swing-like dynamics in which the GSC contributes an effective damping term proportional to active-power deviation, but no explicit virtual inertia term ($\frac{2H}{\omega_0} = 0$) in the sense of the classical swing equation. Instead, additional higher-order dynamics linked to the DC capacitor appear. These properties make ViSynC attractive for robust dc-link voltage regulation [36], but less suited as the primary channel for inertia emulation. In this research we therefore assign complementary tasks to the converters:

- the MSC is used to extract and replenish rotor KE via synthetic inertia and droop control, directly modulating P_{DC} according to frequency signals and rotor speed headroom;
- the GSC operates in ViSynC mode to provide grid-forming synchronisation and dc-link voltage regulation, contributing damping and maintaining current and voltage limits without attempting to emulate additional inertia.

This separation avoids double-counting inertia on the AC side, clarifies the energy path from rotor KE to the grid via DC

capacitor, and reduces interactions between multiple inertia-emulating loops. The proposed framework is shown in Fig. 2.

D. Motivation for PINA-Based Adaptive Control

Given the constraints in (2)–(6) and the heterogeneous operating conditions of WTs, analytically designing gain schedules (K_d, K_p) and recovery trajectories that are near-optimal across all scenarios is difficult. Purely data-driven black-box models can in principle learn such mappings, but they may extrapolate poorly and typically require large labelled datasets.

The proposed approach uses a physics-informed constrained-aware neural adaptive framework to learn maps from $(\Delta f, \dot{f}, \omega_r, \mu)$ to adaptive inertia and droop gains and to a recovery power trajectory, while enforcing feasibility constraints from rotor-energy margins, headroom, and ramp limits, and regularising the learned policy with a swing equation-inspired frequency-response residual to encourage well-shaped transients in the study system. This yields gain schedules that remain interpretable (PD structure), honour converter envelopes by construction (which classic PINN may ignore), and generalise across operating conditions, forming the basis for the GainNet and RecNet controllers introduced in the next section.

The neural networks in PINA are used as compact offline approximators of nonlinear scheduling maps, not as unconstrained replacements for the physical controller. The deployed controller remains an inertia/droop-plus-recovery controller.

A one-dimensional rotor-speed schedule can reflect the available kinetic-energy margin, but it cannot distinguish between a mild event and a severe event at the same rotor speed. Similarly, a rule based only on frequency deviation or RoCoF cannot account for the remaining rotor-energy margin and recovery risk. The desired support action depends jointly on event severity, RoCoF, rotor speed, kinetic-energy headroom, converter limits, and recovery state. This multivariable constrained mapping is difficult to express as a simple analytic schedule without extensive manual tuning.

PINA therefore uses neural networks to learn this mapping offline from simulation data, while preserving the physical controller structure during deployment. In contrast to MPC, no online optimisation is required. In contrast to RL, the network output is not an unconstrained power command but bounded

control parameters that are further processed through physical feasibility constraints.

III. PHYSICS-INFORMED CONSTRAINT-EMBEDDED FREQUENCY CONTROL FRAMEWORK

This section formalises the proposed physics-informed constrained embedded neural adaptive frequency-support controller (GainNet) and its discrete-time implementation. We first describe the pre-filtered signals, then the gain-scheduled control law with hard rails, and finally the physics-informed training objective and deployment pipeline.

A. Signals, Units, and Pre-Filtering

As in (1), Δf and $\dot{f}$ are in p.u. and p.u./s. Rotor speed $\omega(t)$ (pu on turbine base) is constrained to $[\omega_{\min}, \omega_{\max}]$ with typical values $\omega_{\min} \approx 0.7$ p.u. and $\omega_{\max} \approx 1.2$ p.u. To ensure training-deployment parity, identical discrete-time conditioning is applied at sampling period T_s . Each signal is conditioned before it is used by the controller. For a generic input signal x , the deadband (DB) and saturation (SAT) operators are

$$\begin{aligned} \text{DB}(x) &= \text{sign}(x) \max\{|x| - d, 0\}, \\ \text{SAT}(x) &= \min\{\max\{x, -L_{\text{sat}}\}, L_{\text{sat}}\} \end{aligned} \quad (9)$$

where x denotes a generic conditioned signal, applied here to the frequency deviation Δf and its rate $\dot{f}$; $d \geq 0$ is the deadband half-width, and $L_{\text{sat}} > 0$ is the saturation bound. The resulting signal is then filtered by a first-order infinite-impulse-response (IIR) low-pass filter (LPF),

$$y[n] = (1 - \alpha)y[n-1] + \alpha x[n], \quad \alpha = \frac{T_s}{\tau + T_s} \quad (10)$$

RoCoF is obtained by backward difference and LPF,

$$\begin{aligned} \dot{f}_{\text{raw}}[n] &= \frac{\Delta f[n] - \Delta f[n-1]}{T_s}, \\ \dot{f}_{\text{LPF}}[n] &= (1 - \alpha_{\text{roc}})\dot{f}_{\text{LPF}}[n-1] + \alpha_{\text{roc}}\dot{f}_{\text{raw}}[n] \end{aligned} \quad (11)$$

with α_{roc} chosen to balance noise rejection and delay. Rotor kinetic-energy headroom is represented by

$$\mu(\omega) = \text{clip}_{[0,1]} \left(\frac{\omega^2 - \omega_{\min}^2}{\omega_{\max}^2 - \omega_{\min}^2} \right) \quad (12)$$

A disturbance window is detected using the same pre-filtered signals as the controller. A frequency event is triggered when either $|\Delta f_{\text{LPF}}|$ or $|\dot{f}_{\text{LPF}}|$ exceeds a threshold, and exits after both fall below lower thresholds for a dwell time. This produces a binary event mask $m_{\text{evt}}[n] \in \{0, 1\}$ and an event sign

$$\text{sgn}_{\text{evt}}[n] = \text{sign}(\Delta f_{\text{LPF}}[n]) \quad (13)$$

so that under-frequency (UF) and over-frequency (OF) actions can be handled within a single policy.




Fig. 3. Overall runtime flow of the proposed controller. A frequency event triggers GainNet-based adaptive frequency support. After a RoCoF+dwell post-nadir, control transfers to RecNet for energy-aware rotor speed recovery. Upon recovery completion or when Δf subsides, the WT returns to MPPT.

B. Control Law

The external controller adds an auxiliary PD-type frequency-support term to the MPPT branch, as in Sec. II:

$$\Delta P(t) = -k_{\text{in}}(t) \dot{f}(t) - k_{\text{dr}}(t) \Delta f(t) \quad (14)$$

so that the electrical power command becomes

$$P_{\text{cmd}}(t) = P_{\text{MPPT}}(\omega(t)) + \Delta P(t) \quad (15)$$

Because Δf is in p.u. and $\dot{f}$ in p.u./s, the gains map to classical inertia and droop parameters as

$$k_{\text{dr}} = \frac{1}{R}, \quad k_{\text{in}} = 2H \quad (16)$$

with droop R and inertia constant H (s).

Converter and mechanical envelopes are enforced through magnitude and slew-rate limits on the support power,

$$\begin{aligned} \Delta P^+[n] &= \text{SAT}_{\pm \Delta P_{\text{max}}}(\Delta P[n]), \\ |\Delta P^+[n] - \Delta P^+[n-1]| &\leq \gamma T_s \end{aligned} \quad (17)$$

where ΔP_{max} is set by current and torque limits and γ bounds the per-sample ramp rate.

C. Gain Limits and GainNet Architecture

The inertia and droop gains can be calculated according to WT rotor speed and operating status [37] but to remain consistent with NC RfG / GB Grid-Code [38] active power frequency response requirements at the point of connection,

we constrain the adaptive droop gain to the standard LFSM-U/FSM range $R \in [2\%, 12\%]$, i.e.

$$k_{dr} \in [8.33, 50] \quad (\Leftrightarrow R \in [2\%, 12\%]) \quad (18)$$

In the present work these bounds are applied at the (aggregated) turbine level so that the resulting farm-level response remains within the permitted droop range. The inertia-like gain is limited to

$$k_{in} \in [10, 40] \quad (\Leftrightarrow H \in [5, 20] \text{ s}) \quad (19)$$

a design envelope chosen to bracket typical synchronous-generator inertia constants while allowing for stronger synthetic inertia when beneficial.

Both gains are produced by a compact multi-layer perceptron (MLP-GainNet) with hard-cap projection. The GainNet input vector collects the filtered features

$$x = [\Delta f_{LFP}, \dot{f}_{LFP}, \omega, \mu]^T \in \mathbb{R}^4 \quad (20)$$

which are standardised using global train-set statistics,

$$\tilde{x} = \frac{x - \mu_X}{\sigma_X + \varepsilon}, \quad \varepsilon \approx 10^{-6} \quad (21)$$

A three-layer MLP with tanh activations is used,

$$y = g_\theta(\tilde{x}) = W_o \phi(W_3 \phi(W_2 \phi(W_1 \tilde{x} + b_1) + b_2) + b_3) + b_o, \quad \phi(z) = \tanh z \quad (22)$$

and the outputs are mapped to the constrained gains via logistic heads,

$$k_{in} = k_{in,min} + \sigma(y_1)(k_{in,max} - k_{in,min}), \quad k_{dr} = k_{dr,min} + \sigma(y_2)(k_{dr,max} - k_{dr,min}), \quad \sigma(z) = \frac{1}{1 + e^{-z}} \quad (23)$$

When an event is detected ($m_{evt}[n] = 1$), GainNet provides adaptive gains and the resulting $\Delta P[n]$ is applied subject to the hard rails in (17), and the overall runtime flow is summarised in Fig. 3. The recovery network (RecNet) then shapes rotor speed restoration as described in Sec. IV. The power-rotor speed mechanism of WT is shown in Fig. 4.

D. Physics-Informed Objective

In this work, physics information is embedded at two levels. The first is the deployment level: bounded gain heads and command-level rails restrict the controller output to admissible inertia/droop gains, support-power limits, rotor-energy margins, directionality conditions, and slew-rate limits. These constraints are active during online operation and do not depend on the loss function being perfectly minimised. The second is the training level: swing-equation and power-consistency residuals regularise the learned gain surfaces so that the policy remains physically meaningful inside the feasible set. Therefore, residual penalties are used to improve physical consistency, whereas feasibility is enforced by the deployed controller structure.

TABLE II
ROBUSTNESS TO INERTIA/DAMPING ESTIMATION ERRORS: THE INITIAL ESTIMATES (M , D) ARE EACH PERTURBED BY $\pm 50\%$.

M_0	D_0	$\hat{M}$	$\hat{D}$	Nadir (Hz)
-50%	-50%	2.44027	0.44053	49.62916
-50%	+0%	2.44027	0.94053	49.62891
-50%	+50%	2.44027	1.44053	49.62866
+0%	-50%	4.94016	0.44047	49.59748
+0%	+0%	4.94016	0.94047	49.59718
+0%	+50%	4.94016	1.44047	49.59689
+50%	-50%	7.44013	0.44046	49.56124
+50%	+0%	7.44013	0.94046	49.56092
+50%	+50%	7.44013	1.44046	49.56060




Fig. 4. The proposed frequency-support and speed-recovery scheme is illustrated. **Blue (FSP)**: PINA increases power from A to a peak B below $P_{e,max}$, then ramps down to C at P_0 (at the nadir or $\omega_{r,min}$). **Green (SRP)**: RecNet first reduces power along the MPPT curve at the same ω (from C toward D), then restores rotor speed by commanding P_{MPPT} to E. Finally, a shaded-band taper drives the recovery command smoothly to zero.

Let $\Delta P_{sup}[n] = -k_{in}[n]\dot{f}_{LFP}[n] - k_{dr}[n]\Delta f_{LFP}[n]$ denote the discrete support power. GainNet is trained via constraint-embedded learning: the policy is optimised under physics-informed feasibility constraints during detected events. A frequency-response residual is used as a soft regulariser

$$r_{sfr}[n] = M \dot{\Delta f}[n] + D \Delta f[n] - \Delta P_{sup}[n] \quad (24)$$

with $\dot{\Delta f}[n]$ obtained using the same discrete differentiation and LPF structure as in (11). The parameters M and D in (24) are treated as trainable variables and updated jointly with the network weights, preventing the residual from rigidly enforcing a potentially misspecified prior. Their accuracy is, moreover, not critical: feasibility is guaranteed by the bounded output heads, the rails in (17), and the constraints in (26)–(29), none of which depend on M or D . Consequently, inaccuracies in these parameters can only mildly reshape the response within the feasible set. To assess robustness, Table II shows analysis in which the policy is retrained over a (3×3) grid, with the initial values of M and D independently perturbed by $\pm 50\%$, and all nine resulting policies are evaluated on the same nominal plant. The closed-loop frequency nadir varies by at most 0.07 Hz (0.14%), the maximum RoCoF remains unchanged, and the sensitivity to D is approximately two orders of magnitude smaller than that to M . When turbine electrical power and an MPPT reference are available, we




Fig. 5. (a) One-factor-at-a-time sensitivity of the closed-loop frequency nadir to the training-objective weights, each scaled $\times 0.1$ – $\times 10$ about its baseline. (b) log-scale execution-time bars per platform against the $T_s = 1$ ms and 10 ms.

additionally form a measurable consistency residual

$$r_P[n] = (P_e[n] - P_{MPPT}[n]) - \Delta P_{\text{sup}}[n] \quad (25)$$

which directly ties the commanded support to the realised incremental power. The following inequality constraints are enforced during events ($m_{\text{evt}}[n] = 1$). First, the action must oppose the frequency deviation

$$g_{\text{dir}}[n] := \Delta P_{\text{sup}}[n] \Delta f_{\text{LPF}}[n] \leq 0 \quad (26)$$

Second, the support magnitude is bounded by converter and torque limits and rotor-energy margins over short horizons. For UF ($\Delta f_{\text{LPF}} < 0$) we bound positive support by headroom and the instantaneous kinetic-energy extraction rate,

$$\Delta P_{\text{sup}}[n] \leq \min \left\{ \Delta P_{\text{max}}, \text{headroom}[n] + \frac{H_w(\omega[n]^2 - \omega_{\text{min}}^2)}{T_{E,UF}} \right\} \quad (27)$$

and for OF ($\Delta f_{\text{LPF}} > 0$) negative support using the overspeed margin

$$-\Delta P_{\text{sup}}[n] \leq \min \left\{ \Delta P_{\text{max}}, \frac{H_w(\omega_{\text{max}}^2 - \omega[n]^2)}{T_{E,OF}} \right\} \quad (28)$$

where H_w is the turbine inertia constant on the turbine base and $T_{E,UF}, T_{E,OF}$ are energy horizons used to map

TABLE III
COMPUTATIONAL COST AND REAL-TIME EXECUTION OF THE PINA
POLICY FOR $T_s = 10$ MS.

Network complexity and per-step computational cost			
Component	Parameters	MACs	Nonlinearities
GainNet	8,770	8,576	192 tanh
RecNet	1,314	1,248	64 tanh
Total	10,084 (39.4 kB)	9,824	256 tanh + 4 logistic
Execution time on representative platforms			
Platform	Configuration	Time/step	Utilisation
Desktop CPU	C code	2.9 μs	0.03%
TI F28379D*	200 MHz	85–137 μs	0.9–1.4%
ARM Cortex-M7*	480 MHz	35–57 μs	0.4–0.6%
FPGA*	64 MACs, 200 MHz	1–3 μs	0.01–0.03%

*Estimated using 1–2 cycles per MAC and 20–30 cycles per tanh evaluation.




Fig. 6. GainNet training convergence under the objective (32).

the available rotor kinetic-energy margin. Third, we enforce a discrete slew limit consistent with (17):

$$g_{\text{slew}}[n] := \left| \frac{\Delta P_{\text{sup}}[n] - \Delta P_{\text{sup}}[n-1]}{T_s} \right| - \gamma \leq 0 \quad (29)$$

Constraints (26)–(29) are enforced using an augmented Lagrangian (AL) update. For any $g[n] \leq 0$,

$$\begin{aligned} \mathcal{L}_{\text{AL}}(g) &= \lambda^{(k)} \max(0, g) + \frac{\rho}{2} \max(0, g)^2, \\ \lambda^{(k+1)} &= \max(0, \lambda^{(k)} + \rho \mathbb{E}[g]) \end{aligned} \quad (30)$$

booktabs

where ρ increased gradually. To avoid chattering, we penalise gain variation during events:

$$\mathcal{S} = \sum_n m_{\text{evt}}[n] \left(\|k_{\text{in}}[n] - k_{\text{in}}[n-1]\|^2 + \|k_{\text{dr}}[n] - k_{\text{dr}}[n-1]\|^2 \right) \quad (31)$$

The overall GainNet objective becomes

$$\begin{aligned} \mathcal{L}_{\text{gain}} &= \sum_n m_{\text{evt}}[n] \left(w_{\text{sfr}} r_{\text{sfr}}[n]^2 + w_P r_P[n]^2 \right. \\ &\quad \left. + w_{\text{dir}} \mathcal{L}_{\text{AL}}(g_{\text{dir}}[n]) + w_{\text{cap}} \mathcal{L}_{\text{AL}}(g_{\text{cap}}[n]) \right. \\ &\quad \left. + w_{\text{slew}} \mathcal{L}_{\text{AL}}(g_{\text{slew}}[n]) \right) + w_{\text{smooth}} \mathcal{S} \end{aligned} \quad (32)$$

where g_{cap} denotes the appropriate UF/OF cap in (27)–(28). During training, violation rates and multiplier evolution are monitored to confirm compliance with feasibility envelopes. The weights in (32) play heterogeneous roles. The constraint weights scale augmented-Lagrangian terms whose multipliers

adapt according to (30) until the expected violations are driven to zero; converged feasibility is therefore governed by the adaptive multipliers rather than by the specific weight values. The residual weights act as soft physics regularisers that shape the response within the feasible set, whereas the smoothness weight trades gain smoothness against responsiveness.

A one-factor-at-a-time sensitivity study (Fig. 5) retrains the policy with each weighting factor of the training objective scaled over a $(100\times)$ range $((\times 0.1) - (\times 10))$ about its baseline, under an identical seed and evaluation event. Four of the six weights alter the frequency nadir by less than $(0.05, \text{Hz})$ across the entire range, whereas the swing-residual and gain-prior weights exhibit monotone trade-offs, reaching $(0.6, \text{Hz})$ only at the range extremes. All weights remain within $(0.25, \text{Hz})$ for $(\times 1/3) - (\times 3)$ of the baseline value, and the maximum RoCoF remains invariant at $(1.168, \text{Hz/s})$.

Feasibility is preserved throughout the sweep. The hard tier (bounded heads and the rails in (17)) records no violation in any sample of any run, while the soft augmented-Lagrangian constraints exhibit residual exceedances in at most (2.1%) of samples, confined to the first $(\approx 0.4, \text{s})$ of the support transient. Violation rates vary by less than (0.3) percentage points across the sweep, consistent with converged constraint behaviour being governed by the adaptive multipliers rather than the weighting coefficients. The design is therefore insensitive to the precise choice of weights, and the baseline values are reported for reproducibility.

E. Training Protocol and Deployment

Training uses multiple simulated time series sampled at the controller rate, with minibatches formed from time-ordered contiguous segments. The event mask $m_{\text{evt}}[n]$ and sign $\text{sgn}_{\text{evt}}[n]$ are computed from the pre-filtered signals. RoCoF is either logged at the controller rate using the same structure as (11) or recomputed from Δf via (11). Inputs are standardised using global train-set statistics (μ_X, σ_X) computed after filtering. For deployment, the network and preprocessing are realised as a fixed-step discrete-time controller implementation using the exported weights (W_i, b_i) , normalisation statistics, and filter coefficients, ensuring training and deployment parity.

The event-time policy GainNet, together with the swing parameters (M, D) of the residual in (24), is trained offline by backpropagation through the discrete closed-loop dynamics under the objective (32); the recovery policy RecNet is trained separately under the objective (45). Training trajectories are time-contiguous logs sampled at the controller rate $T_s = 10 \text{ ms}$, such that the finite-difference RoCoF of (11), the slew constraint of (17), and the temporal regularisers remain physically meaningful; a disjoint validation split is held out for checkpoint selection. The optimiser is Adam (learning rate 10^{-3} , gradient-norm clipping at 1.0), run for 200 epochs. Feasibility is imposed through the two-tier scheme of Section III-D: the bounded output heads and the hard rails of (17) act on every command, while the augmented-Lagrangian inequality constraints (26)–(29) are enforced through multipliers updated at each epoch with a bounded penalty parameter ρ (capped to prevent penalty runaway). Convergence of the

GainNet objective (32) is shown in Fig. 6. The procedure is insensitive to its hyperparameters: as quantified in Section III-D, scaling any individual loss weight by $\times \frac{1}{3} - \times 3$ changes the closed-loop nadir by at most 0.25 Hz , while even $\times 0.1 - \times 10$ perturbations across the 39 independently retrained policies of the sensitivity and robustness studies leave the hard-tier violation count exactly at zero. Fig. 8 reports the training diagnostics. Panel (a) verifies the training and validation data coverage. Panels (b) and (d) show the learned policy, where the commanded inertia gain increases with rotor speed and event severity. Panel (c) confirms convergence of the physics residual, panel (e) shows that the damping proxy remains positive over the operating domain, and panel (f) shows smooth convergence of the learned gain surface. The computational cost and measured execution time are reported in Table III, while the training configuration is summarised in Table IV.

GainNet employs a layer–neuron architecture of $(4-64-64-64-2)$ with 8,770 trainable parameters, whereas RecNet uses a $(5-32-32-2)$ architecture comprising 1,314 parameters. A single inference of both networks at each control step requires 9,824 multiply–accumulate (MAC) operations, 256 hyperbolic tangent (tanh) evaluations, and four logistic-function evaluations, while storing a total of 10,084 coefficients (39.4 kB in single precision). No online optimisation or matrix factorisation is required. Auto-generated fixed-step C code executes in $2.9 \mu\text{s}$ per control step on a desktop CPU, and Table III reports analytical execution-time estimates based on operation counts and published device throughput for representative embedded targets. The adaptation operates at the supervisory control period, $T_s = 10 \text{ ms}$, such that the worst-case scenario corresponds to less than 1.4% processor utilisation, thereby confirming real-time feasibility.

For the case studies in Section V, we used a single-machine equivalent WPP. Although an aggregated model is employed, the proposed controller is local and agnostic to the wind farm’s internal composition. Each turbine can compute its support from its own measurements; therefore, the farm-level support is given by

$$\Delta P_{\text{farm}}(t) = \sum_{i=1}^N \Delta P_i(\Delta f, \dot{\Delta f}, \omega_i, \mu_i) \quad (33)$$

where the frequency deviation Δf and its rate of change $\dot{\Delta f}$ are common to all turbines, while the rotor speed ω_i and kinetic-energy headroom μ_i are local variables. Parameter and operating-point differences across the farm, including wake-induced variations in incoming wind speed, therefore enter only through the local features (ω_i, μ_i) and are handled by the same controller without modification. Accommodating heterogeneity thus reduces to adapting correctly across the range of operating points spanned by the wind farm. This capability is demonstrated in the time-varying wind case presented in Section V-C, where the same neural networks adapt the inertia and droop gains as wind speed and rotor speed vary significantly, without requiring retuning. The single-machine equivalent adopted for the frequency study represents the coherent aggregate response of the wind farm, while the local nature of the controller ensures that the same support




Fig. 7. Stability of the support loop. (a) Closed-loop eigenvalues over the gain envelope under the lowest-inertia grid condition. (b) Minimum frozen-gain damping ratio across all grid conditions; dashed lines denote the low-inertia band-wise guarantee. (c) Lyapunov-function decay under random gain switching every 10 ms versus the guaranteed bound of (36). (d) Stable inertia-gain ceiling versus RoCoF-filter time constant; the deployed filter lies within the full-envelope stability region, with all remaining behaviour covered by (35).

law remains applicable when individual turbines operate at different operating points.

F. Stability of the Time-Varying Support Loop

This subsection establishes stability for every gain trajectory the controller can command and every switching pattern of the recovery logic. The analysis covers the frequency-support loop in which GainNet and RecNet act; the inner voltage and current loops of the grid-side ViSynC are unchanged from the underlying grid-forming design and are an order of magnitude faster, and the interaction of all loops is validated in the EMT studies of Section V. Each result is stated as an explicit certificate (a norm bound or a set of Lyapunov matrices) whose validity can be verified a posteriori by direct substitution, independently of the solver used to compute it.

1) *Structural property*: Within the admissible set, the support law augments the swing dynamics (24) to

$$(M + k_{\text{in}}(t)) \frac{d\Delta f}{dt} = -(D + k_{\text{dr}}(t)) \Delta f + \Delta p_g - \Delta p_L \quad (34)$$

Since the bounded output heads guarantee $k_{\text{in}}(t), k_{\text{dr}}(t) \geq 0$, the decay rate of (34) is bounded below by $D/(M + k_{\text{in},\text{max}}) > 0$, and $|\Delta f|$ contracts exponentially for arbitrary variation of the gains: by construction, the policy can only ever add inertia and damping.

2) *Unconditional boundedness from the hard tier*: The grid frequency dynamics are open-loop stable, and the deployed command is confined by the magnitude rail of (17). Hence, for any policy (trained or untrained), any gain variation, and any switching,

$$\sup_t |\Delta f(t)| \leq \|h_{u \rightarrow \Delta f}\|_1 \Delta P_{\text{cap}} \leq 0.100 \text{ p.u.} \quad (35)$$

evaluated over the entire $\pm 50\%$ (M, D) polytope of Table II (0.038–0.054 p.u. under the nominal and high-inertia grid conditions; the worst case occurs under the lowest-inertia condition). The hard tier that enforces feasibility (Section III-D) therefore also guarantees that no commandable action can destabilise the frequency dynamics; the learning shapes performance strictly inside this envelope.

3) *Stability guarantee for time-varying gains*: The deployed loop departs from (34) through the primary frequency response, the droop and RoCoF path measurement filters of (11) (time constants $\tau_{\text{dr}} = 0.04$ s and $\tau_{\text{ro}} = 0.10$ s), and the zero-order-hold actuation at $T_s = 10$ ms. With state $x = [\Delta f \ \Delta p_g \ x_{\text{dr}} \ x_{\text{ro}} \ u_{\text{prev}}]^\top$, the rail-inactive closed loop is $x[k+1] = A(\theta_k)x[k]$, with A affine in $\theta = (k_{\text{in}}, k_{\text{dr}})$, so its reachable maps form the convex hull of the four corner matrices of the extended envelope $K_0 = [0, k_{\text{in},\text{max}}] \times [0, k_{\text{dr},\text{max}}] \supset K$, which also contains every command attenuated by the direction gate, the magnitude rail, or the recovery curtailment. For each grid condition of the $\pm 50\%$ (M, D) polytope, Lyapunov matrices satisfying

$$A_v^\top P A_v \preceq (1 - \alpha)P, \quad P \succ 0, \quad \forall A_v, \quad (36)$$

are computed by LMI feasibility and post-verified by direct generalised-eigenvalue substitution; only verified contraction rates are reported.

Under the nominal and high-inertia grid conditions a single P satisfies (36) over the full envelope with $\alpha = 0.0016$ –0.0051 per step (decay factor $\gamma = \sqrt{1 - \alpha} = 0.9975$ –0.9992), so that $V = x^\top P x$ contracts along every trajectory, and the loop is globally uniformly exponentially stable for arbitrary, arbitrarily fast gain variation within K_0 (an envelope-wide guarantee). Under the two low-inertia grid conditions the

guarantee is regime-explicit: the droop gain k_{dr} may vary arbitrarily over its entire range everywhere; the inertia gain k_{in} may vary arbitrarily within any 10 s-wide band of the envelope (per-band $\alpha = 0.0032\text{--}0.0184$) and may traverse its full range at rates up to 1.8–4.0 s/s, equivalently dwelling at least 2.5–5.5 s per band (a band-wise guarantee); still faster inertia-gain variation under the lowest-inertia conditions remains covered by the unconditional bound (35). The band-wise structure has a direct physical origin (with less system inertia there is less filtering of rapid swings of the emulated inertia) and yields a quantitative sampled-data design rule for adaptive inertia emulation. All guarantees are properties of the architecture: the bounded heads define K_0 , so they hold for any trained weights and cannot be invalidated by retraining; the training-time stability regulariser merely biases the learned policy toward the interior of the guaranteed set and carries no deployed role.

4) *Recovery switching and rails*: All RecNet modes command gains within K_0 and share the Lyapunov functions of (36), so stability holds under arbitrary switching among support, recovery, and reattachment, with no dwell-time condition. The recovery curtailment (38) enters as a bounded bias b_k , $|b_k| \leq \Delta P_{rec}$, for which V yields the input-to-state bound

$$\|x[k]\| \leq c_1 \gamma^k \|x[0]\| + c_2 \Delta P_{rec} \quad (37)$$

with constants exported alongside the certificate; the bias vanishes at reattachment, restoring (36). Samples on which the slew rail is active are covered unconditionally by (35); in the constraint studies of Section III-D the rails are never activated by the trained policies, so (36) governs the realised trajectories in their entirety. Rotor-side large-signal protection is enforced by the energy and headroom constraints of Section III-D and is a feasibility property rather than a small-signal one.

5) *Margins and design limits*: Fig. 7 summarises the analysis: the closed-loop eigenvalues over the gain envelope, the minimum frozen-gain damping ratio $\zeta_{min} \geq 0.40$ over the admissible set and all grid conditions, the Lyapunov function decaying under gains switched to a random envelope corner at every 10 ms step against the guaranteed worst-case envelope, and the sampled-data design rule. The stable inertia-gain ceiling falls below the envelope ceiling for RoCoF-filter constants $\tau_{ro} \lesssim 0.09$ s at the 10 ms control rate. At $\tau_{ro} = 0.05$ s, only $k_{in} \leq 20$ s remains stable, because the loop crossover of the inertia path then exceeds the controller Nyquist rate. By contrast, the deployed constants $\tau_{dr} = 0.04$ s, $\tau_{ro} = 0.10$ s, lie inside the region where the entire envelope is stable, with margin. On low-inertia grids the analysis additionally bounds the admissible rate of the inertia gain, tying adaptive-VIC aggressiveness to system inertia quantitatively.

IV. ROTOR SPEED RECOVERY (RECNET)

Following a primary frequency support (PFS) action, rotor kinetic energy is restored by a temporary reduction of electrical power. The commanded active power is

$$P_{cmd}(t) = P_{MPPT}(\omega(t)) + \Delta P_{sup}(t) - \Delta P_{rec}(t) \quad (38)$$

TABLE IV
TRAINING CONFIGURATION

GainNet / RecNet	4-64-64-64-2 / 5-32-32-2 (10,084 coeff.)
Optimiser; learning rate	Adam; 10^{-3} , grad-clip 1.0
Minibatch	full-sequence (contiguous, $T_s=10$ ms)
Validation	disjoint hold-out; best-nadir checkpoint
Feasibility tier	bounded heads + hard rails (17); AL (26)–(29)
AL penalty ρ	10 – 26.5 (capped, monotone)
Residual	swing-eq. (24); $\hat{M}, \hat{D}$ co-identified
GainNet loss weights	$w_s=1$, $w_{sm}=10^{-2}$, $w_{st}=10^{-4}$, $w_{alloc}=10^{-3}$
RecNet loss	schedule-tracking + $\lambda_E, \lambda_H, \lambda_S$ terms
Framework; seeds	PyTorch (CPU); {0, 1, 2}
Inference (deployed)	9,824 MAC/step; 2.9 μ s measured (Table III)

where ΔP_{sup} is the event-time support command from Sec. III, and $\Delta P_{rec} \geq 0$ denotes an active-power curtailment used for rotor speed recovery.

A. Rotor Recovery Activation and Energy Target

Recovery is initiated after the frequency nadir using RoCoF-based finite-state machine with hysteresis and dwell. Let $\dot{f}_{LFP}$ be given by (11), with thresholds $\rho_{enter} > \rho_{exit} > 0$ and dwell time T_{dwell} .

- a frequency event is triggered when Δf crosses threshold.
- the controller enters recovery mode once $\dot{f}_{LFP}$ rises above $-\rho_{exit}$ and persists for at least T_{dwell} .
- if $\dot{f}_{LFP} < -\rho_{enter}$, the controller returns to the event state.

At event to recovery mode transition, a pre-event rotor speed reference ω^* is latched, and the recovery objective is to restore the kinetic energy corresponding to ω^* , and recovery terminates when $\omega \geq \omega^*$. To suppress overshoot and reduce mechanical stress, the recovery action is tapered as ω approaches ω^* using a band $b > 0$

$$\text{taper}(\omega) = \text{clip}_{[0,1]} \left(\frac{\omega^* - \omega}{b} \right) \quad (39)$$

The recovery command is additionally constrained by a magnitude cap u_{cap} (converter/torque limits) and a symmetric rate limiter with ramp-rate limits (R_P^+, R_P^-) .

B. Recovery Power Scheduling

RecNet uses the same filtering and standardisation pipeline as GainNet sec. III. The feature vector is

$$x_{rec} = [\dot{f}_{LFP}, \Delta f_{LFP}, \omega, \mu, \text{headroom}]^T \quad (40)$$

where $\mu(\omega)$ is defined in (12). A compact MLP outputs two bounded recovery parameters through logistic heads: a time constant $\tau \in [\tau_{min}, \tau_{max}]$ and a recovery level $\text{amp} \in [\text{amp}_{min}, \text{amp}_{max}]$.

An internal first-order state s shapes the curtailment smoothly. In recovery mode, s tracks amp with time constant τ ; otherwise s decays to zero with τ_{rel} :

$$\begin{aligned} \text{recovery mode: } s[n+1] &= s[n] + \frac{T_s}{\max(\tau, \epsilon)} (\text{amp} - s[n]), \\ \text{otherwise: } s[n+1] &= s[n] + \frac{T_s}{\max(\tau_{rel}, \epsilon)} (-s[n]) \end{aligned} \quad (41)$$




Fig. 8. GainNet training diagnostics and learned policy: (a) pre-filtered input distributions of a single simulation (Δf_{LPF} , $\dot{f}_{LPF}$, ω , μ) over (train/val) data; (b) k_{in} vs. rotor speed under severe events; (c) validation physics-residual convergence; (d) k_{in} sensitivity to event severity and ω ; (e) damping proxy $(D + k_{dr})/(M + k_{in})$ over the operating domain; (f) evolution of the learned surface $k_{in}(\text{severity}, \omega)$ across epochs.

with $\epsilon > 0$. A short-horizon kinetic-energy prediction is formed using an LPF of $\dot{\omega}$:

$$E_{\text{pred}} = H_w(\omega + \hat{h}\dot{\omega})^2, \quad E_{\text{tgt}} = H_w\omega^{*2} \quad (42)$$

where $\hat{h}$ is a short prediction horizon. The curtailment command prior to rate limiting is $u_{\text{cmd}} = u_{\text{max}} s$ with an energy-to-target cap

$$u_{\text{max}}^E = \max\left\{0, \frac{E_{\text{tgt}} - E_{\text{pred}}}{\max(\tau, \epsilon)}\right\}, \quad (43)$$

$$u_{\text{max}} = \text{taper}(\omega) \cdot \min\{u_{\text{cap}}, u_{\text{max}}^E\}$$

Finally, the applied recovery power is obtained via a symmetric rate limiter,

$$\Delta P_{\text{rec}}[n] = \text{RL}(u_{\text{cmd}}[n]) \quad (44)$$

where $\text{RL}(\cdot)$ enforces the ramp-rate limits (R_P^+ , R_P^-). Recovery ends when $\omega \geq \omega^*$, after which s decays according to (41) and the turbine returns to MPPT.

The recovery design also suppresses secondary frequency dips. Recovery is enabled only after the nadir detector confirms that frequency has passed its minimum, so active-power withdrawal does not begin during the arresting phase. The recovery magnitude is limited by the energy-to-target cap in (43), tapered as the rotor approaches the pre-event speed, and passed through the rate limiter in (44). Therefore, the post-nadir power reduction is gradual and bounded by the remaining rotor-energy deficit, rather than being imposed as an abrupt step. This allows rotor speed to be restored without causing the secondary frequency dip commonly observed in abrupt or timer-based recovery schemes. From the post-nadir power balance, a secondary dip is avoided as long as the recovery withdrawal remains within the growing primary-response reserve margin, i.e. $\Delta P_{\text{rec}}(t) \leq \Delta P_{\text{res}}(t)$ for $t > t_{\text{nadir}}$; the

energy-to-target cap and the rate limit are chosen so that this sufficient condition holds by construction.

C. Training Objective for Recovery Power Shaping

RecNet is trained in a supervised manner using reference targets (τ^*, amp^*) constructed from event severity indicators (e.g., $|\Delta f|$ and $|f|$) and constrained by headroom and rotor-energy budgets. The loss is

$$\mathcal{L}_{\text{rec}} = \sum_n \left(\|(\tau, \text{amp}) - (\tau^*, \text{amp}^*)\|^2 + \lambda_E \max\{0, u_{\text{lb}}[n] - u_{\text{cmd}}[n]\}^2 + \lambda_H \max\{0, u_{\text{cmd}}[n] - u_{\text{cap}}[n]\}^2 + \lambda_S \|(\tau, \text{amp})[n] - (\tau, \text{amp})[n-1]\|^2 \right) \quad (45)$$

so the learned parameters track the reference schedule while penalising violations of the recovery envelope and promoting smooth time evolution.

V. CASE STUDIES AND COMPARISONS

This section evaluates the proposed GainNet and RecNet controllers on the test system shown in Fig. 9 under different wind and load conditions. The system includes two synchronous generators (SGs) and a wind farm, and simulations are carried out in MATLAB® Simulink. SG1 and SG2 are both rated at 100 MVA and are connected to a 60 MW aggregated wind farm, represented by a single equivalent machine comprising four identical 15 MW Type-IV WTGs. The parameters of the WTG system used in this study are given in Table VII. The validation is organised to assess the main claims of the paper through different operating scenarios and benchmark comparisons. These include large-disturbance and low-wind




Fig. 9. The two area power system used for case studies.

cases, time-varying wind operation, different grid strengths, and comparisons with representative control approaches from the literature, including fixed VIC with washout filter, fuzzy AVIC, NN-Blackbox, MPC, RL, and ATFS. The NN-Blackbox controller uses the same structure as the proposed method but is trained without physics-informed constraints or regularisers, thereby isolating the benefit of physics-informed training from purely data-driven learning.

A. Constant High Wind Speed With Large Load Variation

We first consider a high and constant wind speed of $v_w = 12$ m/s. The system operates at the nominal frequency of 50 Hz with the active power load $L_1 = 150$ MW, and all wind turbines initially operate in MPPT mode. At $t = 15$ s, load L_3 is increased by 40 MW, and Fig. 10 illustrates the dynamic response of the system for all control strategies.

The frequency trajectories show that, without frequency support (MPPT case), the load step causes the lowest frequency nadir (around 49.4 Hz) with noticeable oscillations before settling. When VIC or AVIC is enabled to harness KE, the WTs release additional power and the FN improves significantly. Among all frequency support schemes, the proposed PINA controller attains the highest nadir, 49.6 Hz, whereas fixed VIC, NN and fuzzy reach FNs of 49.5 Hz, 49.56 Hz and 49.57 Hz, respectively. The latter three tend to overcompensate, leading to a more oscillatory response around the settling frequency. The dc-link voltage responses further confirm the effectiveness of the proposed technique. The initial sag following the disturbance is most pronounced in the MPPT case, while all frequency support controllers mitigate the initial dip by sharing part of the active-power imbalance. In PINA, dip occurs slightly later due to better KE extraction. During de-loading for rotor speed recovery, the voltage remains slightly lower yet well regulated, indicating that KE-based frequency support does not compromise dc-link stabilisation.

The rotor speed trajectories show how the different controllers utilise the stored kinetic energy. Under MPPT, rotor speed remains unchanged as WTs do not participate in frequency support, whereas all support schemes decelerate the rotor to release KE. The proposed PINA drives the deepest rotor speed reduction, down to 1.03 p.u., indicating that it exploits the available KE more aggressively than the other approaches. RecNet then recovers the rotor speed smoothly back to its nominal value without overshoot or oscillations. In contrast, fixed VIC maintains a higher rotor speed, while NN and Fuzzy utilise more KE than fixed VIC but less than PINA

TABLE V
FREQUENCY SUPPORT METRICS (ΔP_{sup} [p.u.] & $\Delta \omega r_{\text{devi}}$ [p.u.]

Controller	$v_w = 12$ m/s		$v_w = 9$ m/s	
	ΔP_{sup} (p.u)	$\Delta \omega r_{\text{dev}}$ (p.u)	ΔP_{sup} (p.u)	$\Delta \omega r_{\text{dev}}$ (p.u)
Fixed VIC [18]	0.12	0.101	0.04	0.045
NN-Blackbox	0.145	0.140	0.046	0.055
Fuzzy AVIC [15]	0.155	0.142	0.054	0.061
Proposed PINA	0.18	0.170	0.06	0.062




Fig. 10. Simulation results at $v_w = 12$ m/s, with large load disturbance.

and start the recovery phase earlier as they are not explicitly aware of the physical rotor speed safety limits, reflecting under-utilisation of the available KE and frequency support.

The wind farm active power in Fig. 10 demonstrates the trade-off between frequency support and energy capture. Relative to all other controllers, the proposed PINA provides the largest temporary power support $\Delta P_{\text{sup}} = 0.18$ p.u. (2.7 MW per WT) immediately after the disturbance (2.5% more than fuzzy), followed by a controlled de-loading phase to restore the rotor speed. Despite extracting more energy, the duration of the reduced power recovery phase is similar to that of other strategies, so the total energy loss over the event remains limited while the frequency support is enhanced. In comparison, fixed VIC, NN, and fuzzy provide $\Delta P_{\text{sup}} = 0.12$ p.u. (1.8 MW), 0.145 p.u. (2.17 MW), and 0.155 p.u. (2.3 MW), respectively. The results are summarised in Table V. Lastly, the droop and inertia gains generated by GainNet are shown. Fixed VIC keeps both gains constant, while NN and fuzzy vary them with less coordination. In contrast, PINA increases both gains as soon as the frequency drops and gradually reduces them as the rotor decelerates and frequency recovers, consistent with the physical intuition of using high gains when KE margin and event severity are high and reducing them to avoid excessive deceleration and SFDs.




Fig. 11. Simulation results at $v_w = 9$ m/s, with small load disturbance.




Fig. 12. Simulation results for time varying wind speed with random load disturbances.

B. Constant Low Wind Speed With Smaller Load Variation

The second case studies a lower wind resource with $v_w = 9$ m/s. All other initial conditions are kept identical, and a smaller load increase is applied at $t = 15$ s. The corresponding results are presented in Fig. 11. Because the initial rotor speed is lower (around 0.9 p.u.), the available kinetic energy is reduced. This is clearly reflected in the GainNet outputs: the droop and inertia gains in Fig. 11 are lower than in the previous case (decreasing from 30 to 26 and from 40 to 26, respectively), and their post-disturbance variations at 15 s are

also smaller. This indicates that the controller automatically scales down the aggressiveness of the frequency support to keep the rotor speed within a safe operating range. Even under this constrained condition, the proposed PINA controller still achieves the highest frequency nadir (49.72 Hz) and the fastest frequency recovery among all strategies. The improvement over MPPT remains substantial (nadir: 49.47 Hz), while the frequency profile is better and less oscillatory than with fuzzy AVIC (nadir: 49.7 Hz) and the NN-Blackbox controller (nadir: 49.68 Hz). The dc-link voltage dip is smaller than in the previous case due to the smaller disturbance. As before, the proposed controller performed better and maintains a stable dc-link voltage while providing frequency support.

In terms of active power, Fig. 11 shows that the proposed controller delivers a short power surge above the MPPT level followed by a modest de-loading period. The magnitude of the power boost is smaller than in the 12 m/s case, consistent with the lower available kinetic energy. Nevertheless, the temporary boost is still higher than that of the other control strategies $\Delta P_{\text{sup}} = 0.06$ p.u., (0.9 MW per WT). In comparison, the fixed VIC, NN, and fuzzy controllers provide ΔP_{sup} values of 0.04 p.u. (0.6 MW), 0.046 p.u. (0.69 MW), and 0.054 p.u. (0.81 MW), respectively. This shows that the PINA-based controller performs better even at low wind speeds, adapts across wind operating points without retuning, and maintains a favourable compromise between frequency support and energy yield.

C. Time-Varying Wind Speed

The final case investigates the robustness of the proposed controller under time-varying wind speed. The wind speed is varied over a wide range so that the rotor speed spans both under-rated and over-rated regions, while a few load disturbances are also applied randomly over a 100 s simulation window, and Fig. 12 shows the system response.

The frequency trajectories indicate that the proposed PINA provides the best support throughout the simulation, achieving the highest nadir and fastest damping for each event and remaining above fixed VIC, NN-Blackbox, and fuzzy AVIC. This is most apparent when rotor speed is high (e.g., 20s and 70s), where more kinetic energy is available and PINA injects more power to improve the nadir. When wind speed drops and rotor speed approaches its lower limit, the controller moderates its response to prevent rotor stalling while still outperforming the other strategies. The dc-link voltage confirms that the proposed method does not destabilise converter dc dynamics: deviations remain small and quickly return to nominal, with PINA generally closest to 1 p.u.

The rotor speed and wind-farm power profiles in Fig. 12 show how GainNet and RecNet interact under changing operating conditions. PINA exploits high-speed intervals to provide stronger support with larger but controlled speed reductions, followed by smooth RecNet recovery. When wind speed drops, it reduces injection to keep rotor speed above its minimum limit, avoiding protection trips or abrupt curtailment. Rotor speed remains within the admissible interval, with no SFDs. Finally, the droop and inertia gains exhibit piecewise-adaptive behaviour: they are high when frequency deviation and KE




Fig. 13. Frequency support under different grid strengths (SCR = 2, 3, 5) with the trained networks held fixed: (a) system frequency, (b) rotor speed, and (c) wind-farm active power.

margin are large, and step down when wind decreases or frequency recovers. In contrast, fixed VIC uses constant gains, while NN-Blackbox and fuzzy AVIC vary gains without clear dependence on the physical state. This highlights the benefit of embedding energy margins and feasibility constraints, yielding physically meaningful gain scheduling that generalises across varying wind and load patterns.

D. Robustness to Grid Strength

The preceding case studies all use a fixed grid strength. To test whether the proposed controller generalises across grid strengths without retraining, the short-circuit ratio (SCR) at the wind-farm point of connection is varied through the series grid impedance Z_g , with the trained GainNet and RecNet held fixed and no retuning of weights, gains, or constraints. Three grid strengths are examined, SCR = 5 (strong), 3 (moderate), and 2 (weak), under the same wind and load conditions as the low-wind case in Section V-B. The resulting frequency, rotor-speed, and wind-farm power responses are shown in Fig. 13. The controller behaves consistently across all three grid strengths. As the grid weakens, the frequency nadir deepens monotonically, from 49.73 Hz at SCR = 5 to 49.66 Hz at SCR = 2. The settling frequency, by contrast, is essentially identical for all SCR, because the steady-state offset is set by the droop characteristic and the load-generation balance rather than by grid strength. The rotor-speed and wind-farm power trajectories are likewise nearly identical: in every case the turbine releases comparable kinetic energy and follows a similar support-then-recovery profile, differing only marginally on the weakest grid, where the de-loading excursion is slightly smaller and recovery begins slightly earlier. Notably, the loop remains stable and well damped even at SCR = 2, a weak-grid condition under which grid-following designs typically lose stability. Overall, these results confirm that the proposed physics-informed controller generalises from strong to weak grids without retraining, retaining its frequency-support and rotor-recovery behaviour across a wide range of grid strengths.

E. Comparison With Advanced Frequency Support Strategies

The previous comparisons addressed conventional and heuristic designs. To position the proposed framework against the advanced control techniques discussed in Section I, this subsection compares PINA with three representative techniques: (i) an adaptive-gain model predictive controller




Fig. 14. Comparison with advanced control strategies at $v_w = 12$ m/s under a large load disturbance.

TABLE VI
COMPARISON WITH ADVANCED CONTROL PARADIGMS: KEY METRICS

Case	Metric	PINA	MPC [24]	RL [25]	ATFS [23]
12 m/s	Nadir [Hz]	49.61	49.59	49.569	49.578
	Peak WF power [p.u.]	1.10	1.06	1.07	1.07
	$\omega_{r,min}$ [p.u.]	1.040	1.041	1.044	1.037
	$V_{dc,min}$ [p.u.]	0.963	0.961	0.960	0.961
9 m/s	Nadir [Hz]	49.71	49.70	49.68	49.70
	Peak WF power [p.u.]	0.615	0.613	0.61	0.614
	$\omega_{r,min}$ [p.u.]	0.838	0.847	0.843	0.84
	$V_{dc,min}$ [p.u.]	0.98	0.981	0.98	0.979

(MPC) [24], (ii) a reinforcement-learning (RL) controller [25], and (iii) the adaptive temporary frequency support (ATFS) scheme [23]. For a fair comparison, all four controllers run in the same EMT model at the same controller update rate, act on identically filtered frequency and rotor-speed measurements, and are subject to the same support-power magnitude limits and rotor-speed protection rails; only the grid strength and control law generating the support command differs.

Fig. 14 shows the high-wind case ($v_w = 12$ m/s, large load disturbance), and Table VI summarises the metrics. The proposed method attains the highest nadir (49.61 Hz), followed by MPC (49.59 Hz), ATFS (49.58 Hz), and RL (49.57 Hz). The rotor-speed and power traces explain the ordering. PINA injects the largest support peak at the wind-farm terminals

(1.10 p.u. against a 0.93 p.u. pre-event level, versus 1.06–1.07 p.u. for MPC and ATFS) and decelerates the rotor from 1.2 to 1.040 p.u., after which RecNet restores it smoothly and monotonically with no secondary frequency dip. ATFS decelerates the rotor furthest of the set (1.037 p.u.), i.e. it ultimately releases the most kinetic energy, yet records a deeper nadir than the proposed method: its severity- and speed-scaled ceiling bounds the pre-nadir injection by design, so part of the released energy arrives after the nadir, where it no longer improves the minimum frequency. The RL policy shows the opposite signature, it withdraws its support earliest and releases the least kinetic energy (rotor minimum 1.044 p.u.), and correspondingly records the deepest nadir. The comparison therefore illustrates that when rotor energy is released matters as much as the state-adaptive gains concentrate the injection in the arresting phase, where it is most valuable. The dc-link voltage remains well regulated for all techniques, dipping to no less than 0.960 p.u., with the proposed method exhibiting the shallowest excursion (0.963 p.u.) and settling smoothly.

Fig. 15 shows the low-wind case ($v_w = 9$ m/s, small load disturbance), where the available kinetic energy is limited and all methods operate against the same rotor-protection constraints. The peak support powers are nearly identical (0.610–0.615 p.u.), confirming that in a headroom-limited regime the achievable injection is set by the physical energy budget rather than by the control law. Even so, the proposed method again attains the highest nadir (49.71 Hz, within a spread of about 30 mHz over MPC and ATFS at 49.70 Hz and RL at 49.68 Hz) by extracting the most energy within the budget (rotor minimum 0.838 p.u. versus 0.840–0.847 p.u.) and timing its release against the measured frequency state, while its frequency trace recovers earliest after the nadir. The dc-link excursion stays within about 2% of nominal for every method (0.979–0.981 p.u.).

Table VI supports two observations. First, the proposed controller achieves the highest frequency nadir of the set in both cases, by 30–40 mHz under the large disturbance and by 20–30 mHz under the small one; these margins are modest in absolute terms, but they are obtained without sacrificing any of the structural guarantees. Second, the structural properties, rather than the transient metrics, distinguish the methods most clearly. Among the techniques considered, only the proposed controller combines online gain adaptation, feasibility enforced by the controller architecture, a stability certificate over the full commandable gain envelope (Section III-F), interpretable inertia-droop actions, and a deterministic microsecond-scale inference cost. The MPC achieves comparable performance only by solving a constrained optimisation at every control step with an explicit prediction model; the RL policy is a black box whose constraint satisfaction and stability are not guaranteed outside its training distribution; and ATFS follows a predesigned trajectory whose support ceiling is fixed at the event onset. These structural differences, summarised in Table I, are the properties targeted by the physics-informed, constraint-embedded design.




Fig. 15. Comparison with advanced control strategies at $v_w = 9$ m/s under a large load disturbance.

TABLE VII
MAIN PARAMETERS OF WTG SYSTEM

Parameters	Value
Rated power, P_{rated}	15 MVA
Stator resistance, R_s	0.0368 Ω
Stator q -axis inductance, L_q	0.0058 H
Stator d -axis inductance, L_d	0.0087 H
Rated speed, ω_r	0.8 rad/s
Number of pole pairs, p	162
Total reflected inertia, J	3.16×10^8 kg m ²
DC link capacitance	4000 μ F
DC link voltage	10 kV
Output LC filter	0.275 mH; 1024 μ F
Rated wind speed, v_{rated}	12 m/s
Turbine rotor diameter, D	236 m

VI. CONCLUSION

This paper presented a physics-informed constraint-embedded frequency support framework for grid-forming wind turbines that fully exploits rotor kinetic energy while preserving rotor speed safety and dc-link stability. The key idea is a clear separation of roles between the converters: the machine-side converter (MSC) provides synthetic inertia and droop response by modulating rotor-side power, whereas the grid-side converter (GSC) operates in ViSynC mode to regulate the dc-link voltage and grid synchronisation without emulating additional inertia. On top of this structure, a physics-informed constraint-embedded gain network (GainNet) was introduced to adapt inertia and droop gains in real time as functions of frequency deviation, RoCoF, rotor speed, and kinetic energy-headroom, under hard envelopes. A coordinated recovery network (RecNet) shapes a post-nadir, energy-aware rotor speed restoration using energy-to-target capping and rate limits. Case studies show the proposed controller consistently outperforms recent techniques, improving frequency nadir and RoCoF while maintaining dc-link voltage close to nominal and recovering rotor speed without secondary frequency dips under both constant and time-varying wind. Future work will extend the framework to other GFM WT architectures in which the

GSC also emulates virtual inertia (e.g., VSM, dual-port).

REFERENCES

- [1] Z. Zhao, J. Xie, J. Xu, L. Xi, S. Gong, J. Lu, Y. Zeng, Q. Ni, and L. L. Lai, "Assessment and mitigation of multi-mode oscillations in wind-solar hybrid multi-microgrids," *IEEE Transactions on Smart Grid*, vol. 15, no. 2, pp. 1330–1345, 2023.
- [2] G. Zhang, F. Zhang, L. Ding, K. Meng, and Z. Y. Dong, "Wind farm level coordination for optimal inertial control with a second-order cone predictive model," *IEEE Transactions on Sustainable Energy*, vol. 12, no. 4, pp. 2353–2366, 2021.
- [3] NESO, "Technical report on the events of 9 August 2019," National Grid ESO, London, U.K., Tech. Rep., 2019.
- [4] H. Luo, Z. Hu, H. Zhang, and H. Chen, "Coordinated active power control strategy for deloaded wind turbines to improve regulation performance in agc," *IEEE Transactions on Power Systems*, vol. 34, no. 1, pp. 98–108, 2018.
- [5] R. Poudel, V. Krishnan, J. Reilly, P. Koralewicz, and I. Baring-Gould, "Integration of storage in the dc link of a full converter-based distributed wind turbine," in *2021 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE, 2021, pp. 1–5.
- [6] X. Tang, M. Yin, C. Shen, Y. Xu, Z. Y. Dong, and Y. Zou, "Active power control of wind turbine generators via coordinated rotor speed and pitch angle regulation," *IEEE Transactions on Sustainable Energy*, vol. 10, no. 2, pp. 822–832, 2018.
- [7] G. Tu, Y. Li, and J. Xiang, "Coordinated rotor speed and pitch angle control of wind turbines for accurate and efficient frequency response," *IEEE Trans. on Power Systems*, vol. 37, no. 5, pp. 3566–3576, 2021.
- [8] J. Morren, S. W. De Haan, W. L. Kling, and J. Ferreira, "Wind turbines emulating inertia and supporting primary frequency control," *IEEE Transactions on power systems*, vol. 21, no. 1, pp. 433–434, 2006.
- [9] J. Ekanayake and N. Jenkins, "Comparison of the response of doubly fed and fixed-speed induction generator wind turbines to changes in network frequency," *IEEE Transactions on Energy conversion*, vol. 19, no. 4, pp. 800–802, 2004.
- [10] H. Ye, W. Pei, and Z. Qi, "Analytical modeling of inertial and droop responses from a wind farm for short-term frequency regulation in power syst," *IEEE Trans. on Power Syst.*, vol. 31, no. 5, pp. 3414–3423, 2015.
- [11] S. Jiang, H. Wang, G. Li, Y. Xin, L. Wang, and Y. Xu, "A novel coordinated control strategy for frequency regulation of mmc-hvdc connecting offshore wind farms," *IEEE Transactions on Sustainable Energy*, vol. 15, no. 2, pp. 1028–1038, 2023.
- [12] Y. Zhang, S. Yang, Z. Hao, and B. Zhang, "Model-free fast frequency support of wind farms for tracking optimal frequency trajectory," *IEEE Transactions on Sustainable Energy*, 2024.
- [13] J. Lee, E. Muljadi, P. Srensen, and Y. C. Kang, "Releasable kinetic energy-based inertial control of a dfig wind power plant," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 1, pp. 279–288, 2015.
- [14] M. Hwang, E. Muljadi, G. Jang, and Y. C. Kang, "Disturbance-adaptive short-term frequency support of a dfig associated with the variable gain based on the rocof and rotor speed," *IEEE Transactions on Power Systems*, vol. 32, no. 3, pp. 1873–1881, 2016.
- [15] Y. Tang, P. Yang, Y. Yang, Z. Zhao, and L. L. Lai, "Fuzzy adaptive frequency support control strategy for wind turbines with improved rotor speed recovery," *IEEE Transactions on Sustainable Energy*, vol. 15, no. 2, pp. 1351–1364, 2023.
- [16] M. Mehrabankhomartash, M. Saeedifard, and A. Yazdani, "Adjustable wind farm frequency support through multi-terminal hvdc grids," *IEEE Trans. on Sustainable Energy*, vol. 12, no. 2, pp. 1461–1472, 2021.
- [17] Y. Yao, W. Yao, Y. Xiong, S. Lin, H. Zhou, J. Wen, and S. Cheng, "Distributed adaptive leaderless consensus control of wind farm for fast frequency support," *IEEE Transactions on Industrial Electronics*, vol. 71, no. 10, pp. 12 255–12 266, 2024.
- [18] J. Van de Vyver, J. D. De Kooning, B. Meersman, L. Vandeveld, and T. L. Vandoorn, "Droop control as an alternative inertial response strategy for the synthetic inertia on wind turbines," *IEEE Transactions on Power Systems*, vol. 31, no. 2, pp. 1129–1138, 2015.
- [19] K. Liu, Y. Qu, H. Kim, and H. Song, "Avoiding frequency second dip in power unreserved control during wind power rotational speed recovery," *IEEE trans. on power syst.*, vol. 33, no. 3, pp. 3097–3106, 2017.
- [20] X. Peng, W. Yao, C. Yan, J. Wen, and S. Cheng, "Two-stage variable proportion coefficient based frequency support of grid-connected dfi-gws," *IEEE Trans. on Power Systems*, vol. 35, no. 2, pp. 962–974, 2019.
- [21] X. Guo, D. Zhu, J. Hu, X. Zou, Y. Kang, and J. M. Guerrero, "Inertial pll of grid-connected converter for fast frequency support," *CSEE Journal of power and energy Systems*, vol. 9, no. 4, pp. 1594–1599, 2022.
- [22] D. Zhu, X. Guo, B. Tang, J. Hu, X. Zou, and Y. Kang, "Feedforward frequency deviation control in pll for fast inertial response of dfi-g-based wind turbines," *IEEE Transactions on Power Electronics*, vol. 39, no. 1, pp. 664–676, 2023.
- [23] Y. Zhou, D. Zhu, X. Zou, C. He, J. Hu, and Y. Kang, "Adaptive temporary frequency support for dfi-g-based wind turbines," *IEEE Transactions on Energy Conversion*, vol. 38, no. 3, pp. 1937–1949, 2023.
- [24] T. Kerdphol, F. S. Rahman, Y. Mitani, K. Hongesombut, and S. Küfeoglu, "Virtual inertia control-based model predictive control for microgrid frequency stabilization considering high renewable energy integration," *Sustainability*, vol. 9, no. 5, p. 773, 2017.
- [25] Y. Liang, X. Zhao, and L. Sun, "A multiagent reinforcement learning approach for wind farm frequency control," *IEEE Transactions on Industrial Informatics*, vol. 19, no. 2, pp. 1725–1734, 2022.
- [26] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational physics*, vol. 378, pp. 686–707, 2019.
- [27] J. Stiasny, G. S. Misyris, and S. Chatzivasileiadis, "Physics-informed neural networks for non-linear system identification for power system dynamics," in *2021 IEEE Madrid PowerTech*. IEEE, 2021, pp. 1–6.
- [28] R. Nelliakath and S. Chatzivasileiadis, "Physics-informed neural networks for ac optimal power flow," *Electric Power Systems Research*, vol. 212, p. 108412, 2022.
- [29] X. Shi, D. Guo, W. Wang, Y. Cao, M. Shi, Y. Li, and L. Zhou, "Physical-informed detailed aggregation of wind farm based on symbolic regression: Structure, solution and generalizability," *IEEE Transactions on Sustainable Energy*, 2025.
- [30] G. Xu, L. Xu, D. J. Morrow, and D. Chen, "Coordinated dc voltage control of wind turbine with embedded energy storage system," *IEEE transactions on energy conversion*, vol. 27, no. 4, pp. 1036–1045, 2012.
- [31] S. Liu, H. Wu, T. Bosma, and X. Wang, "Impact of dc-link voltage control on torsional vibrations in grid-forming pmsg wind turbines," *IEEE Trans. on Energy Convers.*, vol. 39, no. 4, pp. 2631–2642, 2024.
- [32] L. Chen, X. Du, B. Hu, and F. Blaabjerg, "Drivetrain oscillation analysis of grid forming type-iv wind turbine," *IEEE Transactions on Energy Conversion*, vol. 37, no. 4, pp. 2321–2337, 2022.
- [33] L. Huang, H. Xin, Z. Wang, K. Wu, H. Wang, J. Hu, and C. Lu, "A virtual synchronous control for voltage-source converters utilizing dynamics of dc-link capacitor to realize self-synchronization," *IEEE Journal of Emerging and Selected Topics in Power Electronics*, vol. 5, no. 4, pp. 1565–1577, 2017.
- [34] L. Chen and F. Blaabjerg, "Virtual synchronous generator based on type-iv wind turbine with supercapacitor as storage," in *2021 IEEE/IAS Industrial and Commercial Power System Asia (I&CPS Asia)*. IEEE, 2021, pp. 1194–1200.
- [35] Z. Zhang, D. Sun, C. Zhao, and H. Nian, "An improved grid-forming control strategy of wind turbine generators with the supercapacitor for optimizing primary frequency regulation ability," *IEEE Transactions on Energy Conversion*, 2024.
- [36] J. Girona-Badia, V. A. Lacerda, D. W. Spier, E. Prieto-Araujo, and O. Gomis-Bellmunt, "Resource-aware grid-forming synchronization control: Design, analysis and validation," *IEEE Transactions on Energy Conversion*, 2024.
- [37] M. Toulabi, A. S. Dobakhshari, and A. M. Ranjbar, "An adaptive feedback linearization approach to inertial frequency response of wind turbines," *IEEE Transactions on Sustainable Energy*, vol. 8, no. 3, pp. 916–926, 2016.
- [38] NESO, "Initial RfG Frequency Parameter Selection," https://www.nationalgrid.com/sites/default/files.../RfG_Frequency_Parameters.pdf.