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.

Joint frequency and voltage support from inverter air-conditioner fleets through network-aware predictive control

Divine Okeke

Divine.Okeke@skoltech.ru

Skolkovo Institute of Science and Technology

Elena Gryazina

Skolkovo Institute of Science and Technology

Federico M. Ibanez

Skolkovo Institute of Science and Technology

Yasir Khan

Skolkovo Institute of Science and Technology

Research Article

Keywords: Demand response, Thermostatically controlled loads, Inverter air conditioners, Model predictive control, Frequency regulation, Distribution voltage control

Posted Date: October 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11552393/v1

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

Additional Declarations: No competing interests reported.

Joint frequency and voltage support from inverter air-conditioner fleets through network-aware predictive control

Divine Okeke[1*], Elena Gryazina[1], Federico M. Ibanez[1], Yasir Khan[1]

[1*]Center for Energy Science and Technology, Skolkovo Institute of Science and Technology, Moscow, Russia.

*Corresponding author(s). E-mail(s): Divine.Okeke@skoltech.ru;
Contributing authors: E.Gryazina@skoltech.ru; FM.Ibanez@skoltech.ru;
Yasir.Khan@skoltech.ru;

Abstract

Inverter-based air-conditioner (IAC) fleets can provide frequency support through compressor active-power modulation and voltage support through inverter reactive-power exchange. Coordinating these services requires joint decisions that account for feeder conditions, thermal dynamics, and shared inverter capacity. This paper develops a network-aware fleet model predictive controller (MPC) that combines frequency and thermal prediction with feeder-minimum voltage feedback, updated active/reactive-power sensitivity, and an apparent-power constraint. The resulting convex optimization selects representative per-IAC commands for aggregate fleet dispatch. For the IEEE 13-node test feeder with a 1200 kW/600 kvar load event, the complete 600 s scenario yields 1100 kvar peak support and a minimum feeder voltage of 0.9499 pu, compared with 0.9274 pu without support and 0.9383 pu with active-power MPC plus local Volt-Var control. The proposed controller retains the latter's 59.910 Hz frequency nadir. Reducing inverter rating from 4.0 to 2.6 kVA lowers peak reactive support to 990 kvar and reduces the active-power ceiling by up to 0.92 kW per IAC, identifying the effect of scarce headroom. Across 26 scenarios using National Institute of Standards and Technology measured-data thermal calibration, mean frequency nadir improves from 59.805 to 59.901 Hz and mean minimum fleet-weighted voltage from 0.9399 to 0.9704 pu, with no comfort violations during the tested short events. Nonlinear frequency, transmission–distribution, and converter simulations assess dynamic response and command tracking. The results establish how feeder-informed active/reactive coordination adds voltage support while preserving compressor frequency response, and how inverter rating and disturbance location determine the attainable service.

Keywords: Demand response, Thermostatically controlled loads, Inverter air conditioners, Model predictive control, Frequency regulation, Distribution voltage control

1 Introduction

Coordinated flexible demand offers a way to support power-system balancing while maintaining the service delivered to consumers. Thermal loads provide this flexibility by adjusting electricity consumption over periods in which stored thermal energy sustains the required indoor conditions. Such control has been developed for load following and regulation in systems with variable generation [1, 10]. Delivering a balancing service

also changes power flows in the distribution network that connects the loads to the grid. Ross et al. showed that load-based frequency regulation increases distribution-voltage variation and that the consequences depend on feeder loading and operating conditions [22]. The control problem is therefore to deliver a useful aggregate frequency response while accounting for distribution-voltage needs and the capabilities of the participating equipment.

Air-conditioning fleets offer a direct means of addressing this problem because their electrical and thermal responses occur on different time scales. The room temperature evolves with heat gains, envelope losses, and cooling, allowing short changes in electrical demand within a permitted comfort band. Foundational stochastic load models describe the collective behavior of thermostatically controlled loads (TCLs) [18]; subsequent aggregate models, state estimation, and control methods translate this flexibility into power trajectories for balancing services [19, 21]. Hao et al. express the aggregate power and energy limits through a battery analogy that preserves the role of thermal parameters and temperature bounds [10]. These results connect aggregate electrical flexibility to the thermal constraints that sustain the service delivered to users.

Inverter-based air conditioners (IACs) extend this flexibility through variable-speed compressor operation. Hui et al. developed thermal and electrical models for using continuously adjustable IAC power in frequency regulation [14]. Later work examined the resulting power-system control loop and stability [15], communication-delay effects [25], and predictive control of air-conditioning loads for frequency regulation [26]. This literature provides a basis for controlling the active-power channel while tracking the thermal state. A fleet equipped with controllable grid-side converters can additionally exchange reactive power, creating a second means of influencing the network without obtaining all voltage support from changes in cooling demand.

Distribution-voltage control provides the complementary foundation. Research on photovoltaic inverters has established how reactive-power control can respond to voltage variations and how local control and network-based dispatch use different information [6, 9, 23]. Air-conditioning loads have also been used directly for voltage regulation. Hua et al. coordinated changes in aggregate operating power by resetting temperatures and assigning participation priorities [12]. Their subsequent IAC study used a thermal-battery representation, clustering, and three-phase power-flow calculations to regulate voltage through active-power adjustment in an unbalanced network [13]. Together, these studies show why voltage response must be related to both the network operating point and the usable flexibility of the loads.

Recent IAC research brings active and reactive control into a predictive framework. Chen and Hui developed a two-stage voltage-quality strategy combining photovoltaic operating reserves with MPC of IAC active and reactive power [2]. Chen et al. extended this direction to distributed consensus MPC, with voltage prediction, active/reactive regulation, and heterogeneous comfort requirements; their hardware-in-the-loop experiments demonstrate voltage-control performance [3]. These advances establish IAC active/reactive coordination for voltage regulation. Assigning compressor active power a concurrent frequency-support objective introduces a further design problem: selecting the two services as a joint trajectory within the same inverter capability and thermal dynamics. This study addresses that problem and quantifies how feeder conditions and inverter headroom determine the resulting frequency and voltage response.

This decision has two physical aspects. First, the voltage response to an aggregate change in P/Q depends on feeder conditions, so a representative terminal voltage alone may not reveal the most depressed feeder voltage. Second, reactive exchange and compressor consumption share an apparent-power limit. Underfrequency load reduction can release reactive headroom, whereas restoring compressor power during continued voltage support can use that headroom again. The service combination consequently depends on the operating trajectory, inverter rating, and voltage response to the fleet action.

This paper develops a network-aware fleet MPC that coordinates the two services in a single convex receding-horizon problem. Feeder-minimum voltage feedback and updated P/Q sensitivity guide reactive support, while frequency and thermal predictions determine the compressor trajectory within a shared apparent-power

circle. The coordinator selects representative per-IAC commands and applies their aggregate effect through fixed feeder placement weights. The principal contributions are:

1. A joint frequency–voltage fleet MPC that couples compressor frequency response and feeder-informed reactive support through thermal prediction and the shared inverter P/Q envelope. Its convex formulation incorporates feeder-minimum feedback, a reactive-support floor, and updated voltage sensitivity in the same dispatch decision.
2. A demonstration that coordinated fleet P/Q control improves feeder-local voltage while retaining the active-power frequency benefit, together with a controlled rating ablation that quantifies how scarce inverter headroom reduces reactive support and the feasible active-power ceiling.
3. A reproducible assessment across 26 scenarios using NIST measured-data thermal calibration, supplemented by nonlinear frequency, transmission–distribution, and converter simulations. The combined evidence identifies disturbance location and inverter rating as determinants of simultaneous service performance.

In the primary feeder-stress experiment, the joint controller supplies up to 1100 kvar over the complete trajectory, raises the minimum feeder voltage to 0.9499 pu from 0.9383 pu with active MPC plus local Volt-Var control, and retains the same 59.910 Hz frequency nadir. This result demonstrates the additional voltage service available from coordinating the fleet's two power channels.

2 Modeling and Control Framework

2.1 Fleet Architecture and Power Conventions

The fleet coordinator selects the active power consumed by a representative IAC, P in kW, and its capacitive reactive injection, Q in kvar. For N participating units, the realized per-IAC response is multiplied by N and distributed over prescribed feeder bus-phase weights. Positive reactive support is represented as negative load kvar in the feeder model. Network awareness enters through the current minimum bus-phase voltage and its sensitivity to the aggregate P/Q response. The thermal state and frequency prediction provide the complementary information needed to select compressor power.

At each sampling instant, the coordinator receives the representative state, previous commands, feeder operating point, and frequency-disturbance forecast. It updates the feeder-voltage model and solves for M future P/Q moves. Only the first move is applied. A local implementation layer maps this command to compressor operation and inverter current, including capability clipping and current/DC-link restrictions. This arrangement separates the slow fleet decision from fast converter tracking while maintaining the physical relationship between the two power channels.

Fig. 1 summarizes the information and actuation paths.

2.2 Thermal Dynamics and Inverter Capability

The cooling model relates indoor temperature T_{in} to ambient temperature T_A, internal heat gain Q_{int}, and delivered cooling Q_{cool}:

$$3600 C_{\text{th}} \dot{T}_{in} = \frac{T_A - T_{in}}{R_{\text{th}}} + Q_{\text{int}} - Q_{\text{cool}}. \tag{1}$$

Here, time is in seconds, R_{th} is in °C/kW, and C_{th} is in kWh/°C. The command map enforces power and ramp limits and supplies a temperature- and loading-dependent coefficient of performance (COP). In the sampled simulation, Q_{cool} is the COP evaluated at the feasible compressor command multiplied by realized electrical power after inverter limits; it is held over each 5 s step. This realized power also drives the feeder and frequency response. For prediction, $Q_{\text{cool}} = \text{COP}_{\text{cal}} P_c$, using a fixed calibrated COP and first-order compressor state P_c to retain affine dynamics. This state is advanced separately and fed back to the predictor. Predicted temperature must remain in the comfort interval $[T_{\min}, T_{\max}]$.

The thermal model is fitted to public National Institute of Standards and Technology (NIST) Net-Zero Energy Residential Test Facility (NZERTF) minute-resolution measurements [11, 20]. Cooling intervals are selected

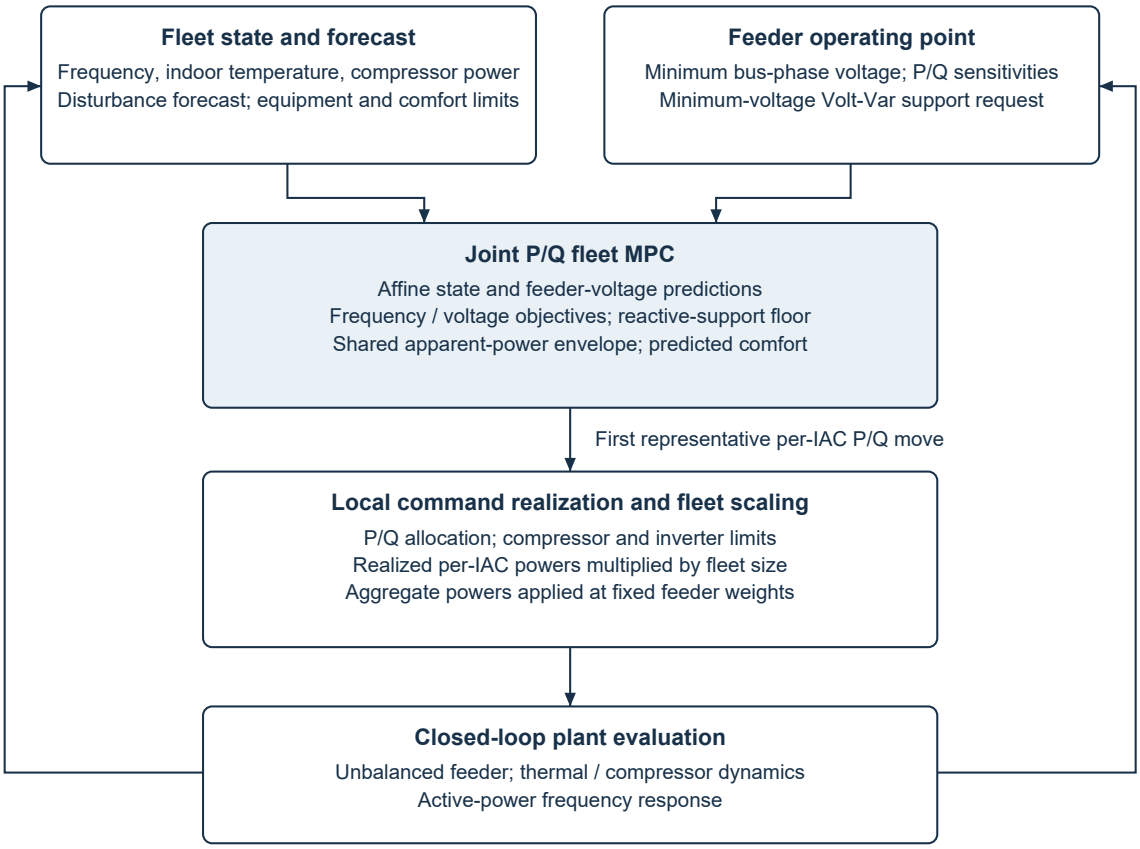


Fig. 1 Network-aware fleet control: feeder-minimum feedback and P/Q sensitivity inform joint optimization; representative commands pass through local device limits before their aggregate response is applied at fixed feeder locations

using heat-pump electrical power, outdoor temperature, and the supply–return air-temperature difference. The fit estimates thermal resistance, capacitance, COP, temperature dependence, and internal gain. It gives an all-window indoor-temperature root-mean-square error (RMSE) of 0.114 °C and a chronological holdout RMSE of 0.173 °C. The fitted parameters represent an effective whole-building model of the instrumented test house; cooling capacity is inferred through COP because it is not directly reported in the minute-resolution heating, ventilation, and air-conditioning (HVAC) file. Calibration plots, holdout and parameter-sensitivity diagnostics, and a separate ResStock/EnergyPlus comparison are provided in Online Resource 1, Section S1.

Compressor and inverter planning limits are

$$P_{\min} \le P \le P_{\max}, \qquad |Q| \le Q_{\max}, \\ \sqrt{P^2 + Q^2} \le S_{\text{rated}}. \tag{2}$$

The apparent-power circle directly couples the two services in the optimizer. Root-mean-square (RMS) current and DC-link limits are implemented in the local command/tracking layers and assessed in representative converter simulations. The local allocation gives reactive support priority when clipping a P/Q request. Compressor ramping is applied after optimization, so the realized power is fed back at the next control step.

2.3 Prediction of Frequency, Temperature, and Feeder Voltage

The representative state is $x = [\Delta f, T_{in}, P_c]^T$, comprising frequency deviation, indoor temperature, and compressor power. Exact zero-order-hold discretization of the reduced continuous model gives

$$x_{k+j+1} = A_d x_{k+j} + B_d P_{k+j} + g_d + E_d d_{k+j}, \quad (3)$$

where d is the exogenous frequency-disturbance input. The frequency equation relates load relief to frequency response, the thermal equation links cooling power to indoor temperature, and the compressor equation represents its first-order response. Reactive power enters the feeder-voltage prediction separately.

Let V_0 be the minimum bus-phase voltage at the current feeder operating point, and let (P_0, Q_0) be the corresponding realized per-IAC powers. Finite differences from unbalanced OpenDSS power flow [7, 8] yield sensitivities s_P and s_Q to aggregate active consumption and reactive injection. The prediction is

$$\hat{V}_{k+j} = V_0 + s_P N(P_{k+j} - P_0) + s_Q N(Q_{k+j} - Q_0). \quad (4)$$

The perturbations use the same fixed placement weights as the applied fleet response; their magnitudes are $\max(25, 0.05N)$ kW and kvar, respectively. The scalar minimum is recomputed for each perturbed power flow and can occur at a different bus-phase. The prediction therefore follows the scalar feeder-minimum metric as its location changes.

The full P/Q controller also uses feeder-minimum feedback to set a capacitive-support floor:

$$Q_{\text{floor}} = \max\{q_{\text{VV}}(V_0), 0\}, \quad (5)$$

where q_{VV} is the piecewise Volt-Var curve defined in Section 3. The floor and the sensitivity-based voltage objective jointly guide reactive support. Both are updated from the present feeder operating point and held over the prediction horizon.

2.4 Joint Receding-Horizon Optimization

The coordinator minimizes

$$\begin{aligned} J = \sum_{j=0}^{M-1} \big[& w_f(\Delta f_{k+j+1} - \Delta f_{\text{ref}})^2 \\ & + w_T(T_{in,k+j+1} - T_{\text{set}})^2 \\ & + w_V \max(0, V_{\text{ref}} - \hat{V}_{k+j})^2 \\ & + w_{\Delta P}\Delta P_{k+j}^2 + w_{\Delta Q}\Delta Q_{k+j}^2 \\ & + w_P(P_{k+j} - P_{\text{set}})^2 + w_Q Q_{k+j}^2 \big], \end{aligned} \quad (6)$$

subject to the measured initial state, dynamics (3), capability constraints (2), and

$$\begin{aligned} & T_{\min} \le T_{in,k+j+1} \le T_{\max}, \\ & V_{\min} \le \hat{V}_{k+j} \le V_{\max}, \qquad Q_{k+j} \ge Q_{\text{floor}}. \end{aligned} \quad (7)$$

The control increments are relative to the preceding horizon move, with the previous applied command used at $j = 0$. The frequency term encourages compressor load relief during underfrequency. Temperature and power tracking encourage recovery toward normal cooling operation; the voltage term penalizes shortfall, while move and reactive-effort penalties regularize the response. All scenarios use $w_f = 60$, $w_T = 3$, $w_V = 500$, $w_{\Delta P} = 0.04$, $w_{\Delta Q} = 0.02$, $w_P = 0.02$, and $w_Q = 0.01$. The voltage reference is 0.96 pu and the hard predicted-voltage bounds are 0.88–1.08 pu. These settings are distinct from the 0.95 pu reference level used to assess voltage recovery in the figures.

The affine dynamics, convex objective, and second-order-cone capability constraint form a convex second-order cone program (SOCP). The implementation uses CVXPY with CLARABEL and warm starts, a 5 s sampling interval, and $M = 12$ steps. The active-only formulation is a quadratic program (QP) solved with OSQP. A solver status other than optimal or optimal-inaccurate triggers a clipped temperature-aware active command and retention of the previous reactive command before local clipping; this status fallback was not exercised in the reported scenarios.

The coupling can be seen directly from the reactive headroom,

$$|Q| \leq \sqrt{S_{\text{rated}}^2 - P^2}. \tag{8}$$

Reducing compressor power can simultaneously support frequency and release kvar capacity. As compressor power recovers, continued reactive support reduces the upper feasible active-power level. Combining this envelope with (4) lets the controller account for both the voltage value and the capacity requirement of its joint action.

3 Evaluation Design

3.1 Feeder, Frequency Event, and Controller Comparisons

The distribution study uses the unbalanced IEEE 13-node feeder in OpenDSS [16, 17]. IAC capacity is attached to the original load buses and phases in proportion to benchmark active loads; a weak-bus-prioritized placement is included as a sensitivity case. The nominal fleet comprises 500 IACs, each with 2.0 kW set power, a 0.6–3.5 kW compressor range, a 4.0 kVA inverter, and a 2.2 kvar reactive limit. The nominal comfort band is ±1 °C around the calibrated 23.42 °C setpoint. Online Resource 1, Section S1 gives the nominal parameters.

A 1200 kW bulk power imbalance is applied from 120 to 420 s. Closed-loop frequency is obtained from a dynamic replay model with inertia, damping, governor droop, and turbine-governor lag; its load-relief response kernel maps the realized fleet active-power deviation into frequency response. The primary feeder-local stress adds 1200 kW/600 kvar over the original feeder load locations from 180 to 380 s, then retains half this increment (600 kW/300 kvar) until 460 s before removal. The 600/300 and 900/450 kW/kvar cases use the same two-stage profile at lower amplitudes. A separate source-sag family reduces upstream voltage, including a 0.92 pu event from 180 to 380 s, followed by recovery to 0.985 pu until 460 s and then to 1.0 pu. These cases distinguish voltage changes originating within the feeder from those imposed at its source.

Five controllers use common device ratings and downstream command-processing layers:

- No support maintains the compressor set power and zero kvar.
- Rule based combines Frequency-Watt response with local Volt-Var droop.
- MPC-P uses active-power MPC with reactive support disabled.
- MPC-P+VV combines active-power MPC with the separate local Volt-Var curve; the active optimization accounts for the requested kvar through its feasible active-power bounds.
- Full P/Q MPC uses the joint formulation in Section 2, including feeder-minimum sensitivity, the feeder-informed reactive floor, and the shared capability circle.

The local Volt-Var input is the representative fleet-weighted point-of-common-coupling (PCC) voltage returned by the feeder model. Its deadband is 0.98–1.02 pu, with linear ramps to full capacitive injection at 0.94 pu and full absorption at 1.06 pu. The Frequency-Watt baseline has a ±0.03 Hz deadband and reaches full response at a 0.30 Hz deviation, with availability scaled across the full comfort band: cooling curtailment decreases toward the hot bound, while additional cooling decreases toward the cold bound. At or beyond the hot and cold bounds, maximum and minimum cooling commands, respectively, override frequency support. The explicit rule is given in Online Resource 1, Section S1. MPC-P+VV and full P/Q MPC therefore compare two complete service architectures: active prediction plus local droop, and joint P/Q prediction with feeder information.

All MPC cases use the prescribed disturbance schedule as their horizon forecast.

The experiment matrix contains 26 scenarios and five controller runs per scenario. It covers source-sag and feeder-stress magnitudes, fleet size, TCL share, placement, comfort band, and inverter headroom, plus three stochastic diagnostic cases. The diagnostics vary selected noise, delay, compressor-time, fleet-size, and sag settings; the fitted NIST thermal parameters remain fixed. The reported means summarize this engineering test suite.

Frequency nadir, minimum feeder bus-phase voltage, minimum fleet-weighted PCC voltage, and peak aggregate kvar are evaluated over the complete 0–600 s trajectory, including recovery.

Feeder minimum and fleet-weighted PCC voltage are separate metrics. Event-window values are identified explicitly when used. Comfort is measured by violation count, violation degree-minutes, and maximum setpoint deviation; computation time is recorded around sensitivity estimation and controller solution.

3.2 Dynamic and Converter Evaluation

Coordinated representative traces are examined with complementary simulation models. ANDES [5] evaluates active-power frequency response on the IEEE 14-bus benchmark with five GENROU machines, TGOV1 governors, exciters, and bus-frequency measurements. Aggregate IAC load relief is mapped to timed governor auxiliary-power alterations on a 100 MVA base. This provides a nonlinear frequency comparison alongside the replay model.

The transmission–distribution (T&D) check exchanges interface voltage and feeder boundary load at 100 ms macro-steps. ANDES supplies the feeder source voltage; OpenDSS evaluates the unbalanced feeder and returns the boundary active-power change for the next transmission-dynamic segment. This explicit loose coupling retains feedback between the domains, with the time-step considerations inherent to such co-simulation [24].

Averaged and switching converter models assess tracking of representative full-P/Q references, including source-sag-family inputs. The averaged model includes phase-locked-loop (PLL) synchronization, dq current control, an L filter, pulse-width modulation (PWM)/computation delay, DC-link capacitance, current limiting, and anti-windup. The switching model combines fundamental dq dynamics with explicit PWM ripple reconstruction for a representative converter at a Thevenin-equivalent feeder bus, 10 kHz bipolar full-bridge PWM, and a 20 μs simulation step; a short 0.75 pu notch and frequency pulse test its fast response. Single-phase synchronization follows established control principles [4]. The coordinated representative inputs connect the fleet-control study to frequency dynamics, boundary feedback, and current tracking at their respective time scales. Python/Matplotlib scripts generate the plots directly from archived simulation outputs, with comfort bounds taken from the calibrated model parameters. The control schematic was prepared with ReportLab.

4 Results

4.1 Simultaneous Support During Feeder-Local Stress

Joint P/Q control adds feeder-voltage support while retaining the active-power frequency response in the primary case. Fig. 2 shows the response to the 1200 kW/600 kvar feeder-load increase, superimposed on the bulk frequency event. Over the complete trajectory, full P/Q MPC raises the minimum feeder voltage from 0.9274 pu without support to 0.9499 pu and retains a 59.910 Hz frequency nadir. Table 1 separates the active-power benefit from the additional voltage benefit: MPC-P and MPC-P+VV reach the same frequency nadir, but their minimum feeder voltage is 0.9383 pu. Full P/Q MPC adds approximately 0.0116 pu to that minimum through the combined feeder-informed P/Q architecture.

Feeder-minimum feedback enables reactive support during a voltage depression that is not visible in the representative PCC metric. In MPC-P+VV, the representative fleet-weighted PCC voltage ranges from 0.9838 to 1.0013 pu and remains inside the local Volt-Var deadband. Its reactive output is therefore zero. Full P/Q MPC uses the feeder-minimum voltage for its support floor and sensitivity prediction and supplies up to 1100 kvar across the trajectory. This activates reactive support where the representative PCC metric does not request it. The neighboring 600/300 and 900/450 kW/kvar stress cases show the same qualitative behavior.

During the full-amplitude interval (180–380 s), full P/Q MPC maintains a minimum feeder voltage of 0.9570 pu and supplies up to 1071 kvar; MPC-P+VV reaches 0.9402 pu. The full-trajectory minimum of 0.9499 pu in Table 1 occurs at 445 s, during the residual 600 kW/300 kvar loading, and remains slightly below 0.95 pu. Evaluating the complete trajectory thus captures the voltage requirement beyond the full-amplitude interval. The autonomous rule first curtails load at 125 s, when the initial nadir is already

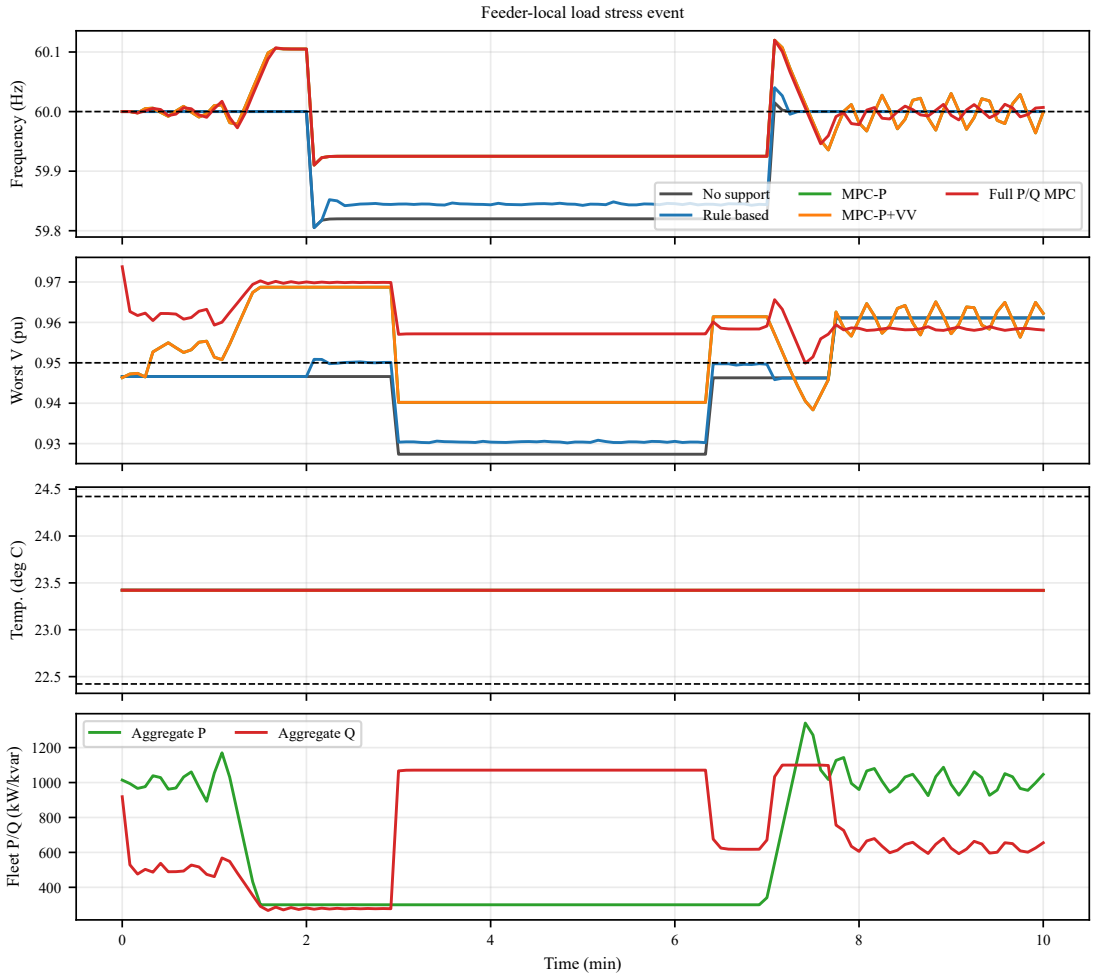


Fig. 2 Frequency and voltage response to the primary feeder-local load increase: full amplitude from 180 to 380 s, half amplitude until 460 s, and zero thereafter. The bottom panel shows aggregate realized full-P/Q active and reactive power. Reported extrema in Table 1 cover the complete 0–600 s trajectory, including recovery

Table 1 Primary 1200 kW/600 kvar feeder-stress case: full-trajectory extrema over 0–600 s. $V_{\min}$ is the minimum feeder bus-phase voltage, and Q_{peak} is peak aggregate capacitive support.

Controller	Q_{peak} (kvar)	$V_{\min}$ (pu)	$f_{\min}$ (Hz)
No support	0	0.9274	59.805
Rule based	0	0.9302	59.805
MPC-P	0	0.9383	59.910
MPC-P+VV	0	0.9383	59.910
Full P/Q MPC	1100	0.9499	59.910

recorded, and then supports frequency recovery. The MPC controllers can act on the scheduled disturbance forecast; their nadir comparison with the measurement-driven rule therefore includes the benefit of preview.

4.2 How Inverter Headroom Shapes the Joint Service

The apparent-power ablation repeats the primary feeder-stress case with 2.6, 4.0, and 6.0 kVA

per IAC while retaining the disturbance, thermal calibration, kvar limit, controller weights, and sensitivity update. For a realized reactive command, the remaining active-power ceiling is

$$P_{\text{ceiling}}(Q) = \min\left(P_{\max}, \sqrt{S_{\text{rated}}^2 - Q^2}\right). \quad (9)$$

Its loss is measured relative to $\min(P_{\max}, S_{\text{rated}})$, the no-kvar ceiling for the same rating. The ceiling quantifies the feasible active-power range available alongside the reactive service.

The 2.6 kVA case reaches a utilization $\sqrt{P^2 + Q^2}/S_{\text{rated}} \geq 0.99$ at 10 samples, classified as near-binding. Its maximum active-ceiling loss is 0.916 kW per IAC, peak aggregate support falls to 990 kvar, and minimum feeder voltage falls to 0.9452 pu. At the nominal 4.0 kVA rating, peak utilization is 0.867 and no samples meet the near-binding criterion. A theoretical ceiling loss of 0.159 kW per IAC still occurs, but the chosen trajectory remains inside the capability circle. At 6.0 kVA, the ceiling loss is zero. Fig. 3 displays these trajectories; Online Resource 1, Section S3 provides the numerical table.

Frequency nadir is essentially unchanged across these ratings because compressor power reaches its lower bound during the main frequency disturbance. Active-load relief therefore supports frequency and releases reactive headroom simultaneously. The rating determines how much reactive support can be sustained and how much active-power capacity remains available over the wider trajectory. In the nominal 4.0 kVA case, the selected joint response is accommodated with spare apparent-power capacity.

4.3 Disturbance Location and Voltage Recovery

Disturbance location determines the voltage improvement attainable from fleet reactive support. With source voltage reduced to 0.92 pu, full P/Q MPC raises the minimum fleet-weighted PCC voltage from 0.9396 to 0.9716 pu, while the minimum feeder bus-phase voltage rises from 0.8988 to 0.9199 pu. Reactive support reduces feeder voltage drops, but the depressed source and unbalanced feeder still constrain the attainable minimum. The source-sag trajectory in Online Resource 1, Section S2 separates the feeder-drop reduction provided by IACs from the source-voltage restoration required upstream.

Increasing inverter kVA, kvar, RMS-current, and DC-link-current headroom together raises minimum fleet-weighted PCC voltage to 0.9847 pu at a 1.5 pu rating in the representative source-sag experiment. The worst feeder voltage remains limited by the upstream disturbance. Additional local capacity thus improves the voltage experienced by the fleet while upstream restoration and feeder voltage-control equipment remain relevant to the most depressed nodes.

A local optimal-power-flow (OPF) sensitivity diagnostic compares affine voltage predictions with nonlinear feeder power flow after applying the corresponding P/Q setpoint. Prediction error is below 0.001 pu for the feeder-local stress case and approximately 0.017 pu for the source-sag case (Online Resource 1, Section S3). This supports the use of updated local sensitivities for command selection together with nonlinear power-flow feedback, particularly under large source disturbances.

4.4 Performance Across the Scenario Set

Across the 26 scenarios, full P/Q MPC improves the mean frequency nadir from 59.805 Hz without support to 59.901 Hz. The mean minimum fleet-weighted PCC voltage increases from 0.9399 to 0.9704 pu, and the mean minimum feeder bus-phase voltage increases from 0.9001 to 0.9215 pu. The distinct voltage averages reflect the different locations represented by the two metrics and the mixture of source-sag and feeder-load events. Online Resource 1, Section S2 provides all-controller summaries and representative-case values. With 40% TCL participation, the frequency nadir reaches approximately 59.93 Hz, illustrating the value of increasing the available active-load response.

All 130 controller–scenario runs retain zero comfort violations and zero violation degree-minutes. The largest absolute departure from the calibrated setpoint is approximately 0.0039 °C, within even the ±0.5 °C sensitivity band. The events are short relative to the effective thermal time constant, so comfort is preserved with substantial margin and does not determine the response limit in this suite.

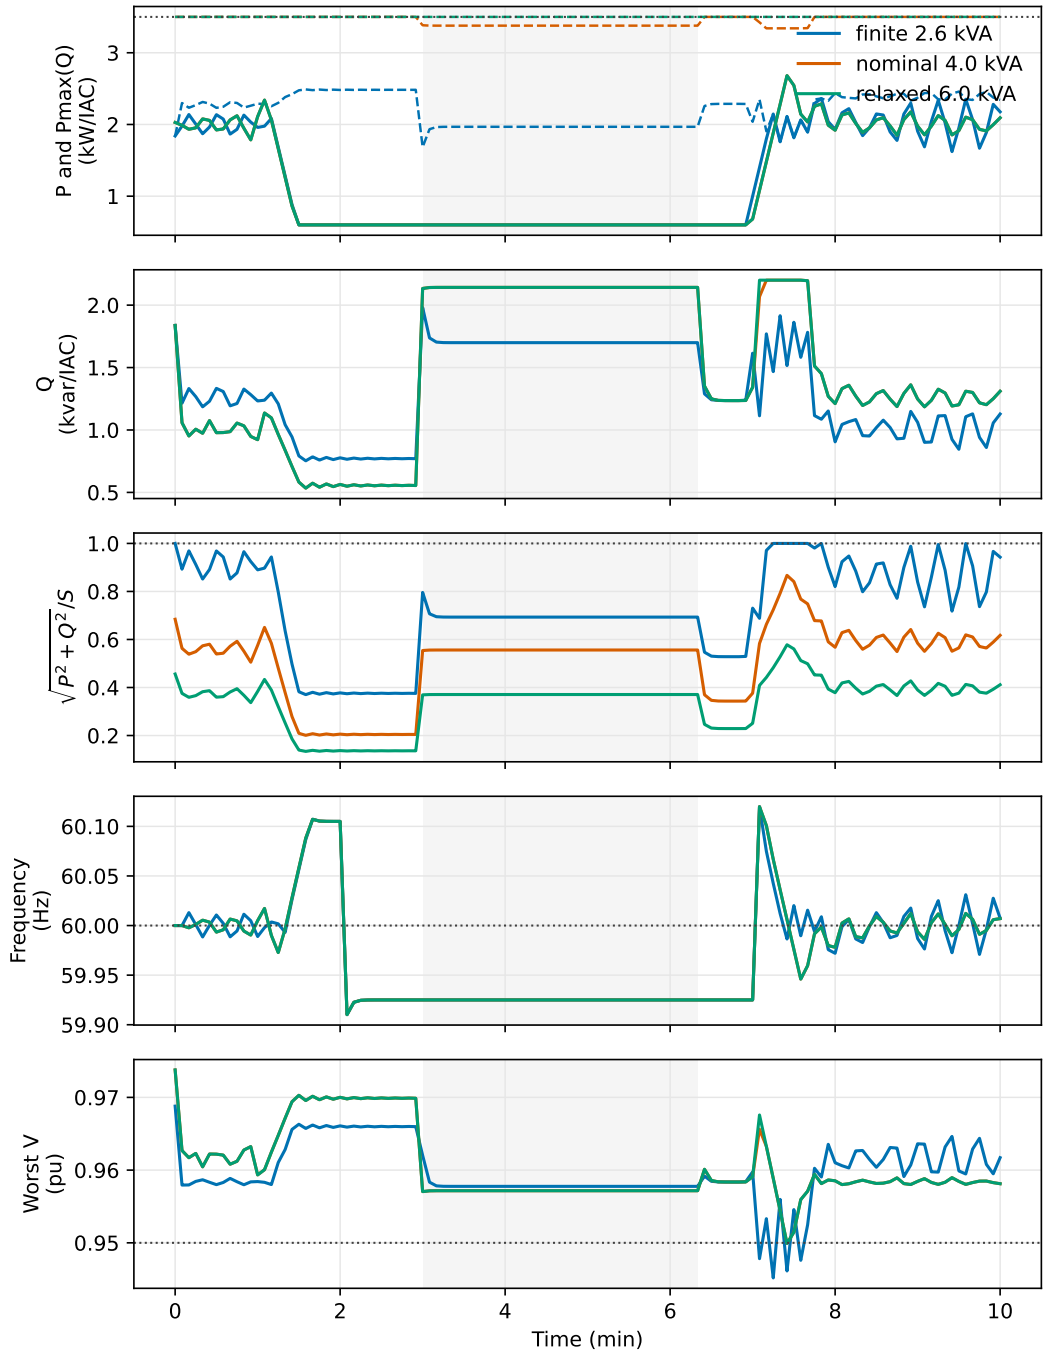


Fig. 3 Apparent-power rating ablation for the primary feeder-stress case. Solid active-power curves show realized per-IAC power; dashed curves show the ceiling left by the simultaneous reactive command. Utilization is the fraction of rated apparent power used. The shaded interval identifies the full-amplitude feeder-load event

For full P/Q MPC, the mean of scenario-mean controller times is 539 ms, the largest scenario mean is 555 ms, and the largest individual recorded sample is 909 ms. These measurements include feeder-sensitivity estimation and optimization and are below the 5 s control interval. The timed block excludes downstream feeder application, plant advancement, and converter simulation.

4.5 Dynamic Response and Converter Tracking

In the nonlinear ANDES check, active-power relief improves the center-of-inertia frequency nadir from 59.800 Hz without support to 59.884 Hz for both MPC-P and full P/Q MPC. The two controllers supply the same active response in this frequency-only benchmark. The frequency improvement is therefore retained under nonlinear machine and governor dynamics.

The explicit T&D source-sag check gives a no-support nadir of 59.826 Hz and a minimum feeder voltage of 0.8988 pu. MPC-P reaches 59.904 Hz. Full P/Q MPC reaches 59.892 Hz while raising the minimum feeder voltage to 0.9185 pu and the minimum fleet-weighted PCC voltage to 0.9700 pu. Joint support therefore retains a frequency improvement over no support in this coupled case, with a nadir 0.012 Hz below active-only MPC. This difference describes the combined closed-loop response, including feeder boundary feedback.

The averaged converter model tracks its representative full-P/Q reference with RMS errors of 0.367 kW and 0.117 kvar, maximum RMS current of 13.31 A against an 18 A limit, no current saturation, and a minimum DC-link voltage of 388.6 V. The switching model gives RMS fundamental tracking errors of 0.012 kW and 0.010 kvar, maximum fundamental RMS current of 4.72 A, no voltage saturation, and a minimum DC-link voltage of 389.2 V. Detailed current, voltage, and PLL traces are reported in Online Resource 1, Section S4. These results demonstrate tracking of the fleet references within the modeled current and DC-link operating conditions.

5 Discussion

The results establish the value of treating frequency and feeder-voltage support as one fleet-control problem. In the primary experiment, joint P/Q MPC retains the frequency nadir achieved by active MPC plus local Volt-Var control and raises the feeder minimum by approximately 0.0116 pu. The combined use of feeder-minimum feedback, a reactive-support floor, sensitivity prediction, and joint optimization supplies kvar while the representative PCC voltage remains in its deadband. The demonstrated benefit belongs to this integrated control architecture.

Two physical factors govern the attainable service. Inverter rating determines the feasible P/Q combinations: nominal spare capacity accommodates both services, while the 2.6 kVA case exposes the reduction in reactive output and active-power ceiling under scarce headroom. The operating trajectory also matters. Underfrequency load reduction releases converter capacity, making frequency and voltage actions complementary during that part of the event; compressor recovery can subsequently use the released headroom. Disturbance location determines the voltage value of the response. Local kvar reduces feeder drops effectively during feeder loading and improves fleet voltage during source sag, while restoration of the upstream voltage remains a separate system requirement. These findings support joint coordination of IAC fleets with feeder and bulk-system controls.

The aggregate formulation makes repeated optimization tractable through representative thermal/electrical states, fixed placement weights, and a scalar minimum-voltage prediction. The 909 ms largest recorded controller sample, including sensitivity estimation, fits within the 5 s control interval in the tested implementation. Updating the sensitivities and evaluating the response with nonlinear feeder power flow accommodate changes in operating point. Spatially resolved extensions could use the same P/Q coordination principle to select individual phases, locations, and heterogeneous device headroom.

The evidence covers short disturbances on the tested IEEE benchmarks with an effective thermal model calibrated to NIST test-house measurements. Comfort constraints retain substantial margin in these events; the calibration window

also identifies thermal capacitance less strongly than cooling and internal gains. Sustained service should therefore be assessed with longer recovery cycles and heterogeneous occupied buildings. The 26-scenario means characterize the tested conditions, and the three stochastic cases provide diagnostic rather than population-level evidence. The quasi-static feeder, explicit 100 ms T&D coupling, and representative converter models resolve complementary aspects of the response. Laboratory and field evaluation would extend this evidence to equipment behavior, communications, and end-to-end operation.

6 Conclusion

This paper developed a network-aware fleet MPC for simultaneous frequency and distribution-voltage support from inverter-based air conditioners. The controller combines frequency and thermal prediction with feeder-minimum voltage feedback, updated P/Q sensitivity, and a shared apparent-power envelope to select joint compressor and reactive-power commands.

The primary 1200 kW/600 kvar feeder-stress experiment demonstrates the resulting service: the full-trajectory minimum feeder voltage rises from 0.9274 pu without support and 0.9383 pu with MPC-P+VV to 0.9499 pu, while full P/Q MPC retains the 59.910 Hz nadir achieved by active-power MPC. Across 26 scenarios, frequency and fleet-weighted voltage improve with no comfort violations during the tested short events. Nonlinear frequency, T&D, and converter simulations further support the dynamic response and tracking of representative control actions.

The experiments also identify the physical conditions that determine performance. Nominal 4.0 kVA capacity accommodates the selected joint trajectory, whereas a 2.6 kVA rating reduces reactive support and the available active-power ceiling. Source-sag results distinguish feeder-drop reduction from upstream voltage restoration. Feeder-informed P/Q coordination therefore provides a control basis for obtaining voltage support alongside compressor frequency response, with inverter sizing, disturbance location, and thermal service duration guiding its extension to spatially resolved and experimentally validated fleet operation.

Supplementary Information

Online Resource 1 contains model parameters, measured-data calibration and sensitivity diagnostics, complete controller comparisons, inverter-headroom and optimal-power-flow results, and supporting frequency, transmission–distribution, and converter simulations (Sections S1–S4). Online Resource 2 contains the executable Python/OpenDSS model, controller definitions, scenario inputs, archived simulation results, figure-generation scripts, and tests supporting the reported comparisons.

Statements and Declarations

Data and Code Availability

The code, scenario definitions, calibration inputs, and simulation results supporting this study are provided in Online Resource 2. The original NIST NZERTF measurements are available from the cited NIST data repository.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT Authorship Contribution Statement

Divine Okeke: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Visualization, Writing - original draft, Writing - review and editing. Elena Gryazina: Conceptualization, Methodology, Supervision, Project administration, Writing - review and editing. Federico M. Ibanez: Methodology, Supervision, Validation, Writing - review and editing. Yasir Khan: Writing - review and editing.

Declaration of Generative AI and AI-Assisted Technologies

During the preparation of this manuscript, the authors used OpenAI ChatGPT for language editing. The authors take full responsibility for the content of the manuscript.

References

[1] Callaway DS (2009) Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. Energy Conversion and Management 50(5):1389–1400. https://doi.org/10.1016/j.enconman.2008.12.012

[2] Chen L, Hui H (2025) Model predictive control-based active/reactive power regulation of inverter air conditioners for improving voltage quality of distribution systems. IEEE Transactions on Industrial Informatics 21(1):922–931. https://doi.org/10.1109/TII.2024.3468475

[3] Chen L, Hui H, Song Y, et al (2026) Distributed consensus-based model predictive control of inverter air conditioners for voltage regulation in distribution systems. IEEE Transactions on Power Systems 41(4):2441–2453. https://doi.org/10.1109/TPWRS.2026.3665716

[4] Ciobotaru M, Teodorescu R, Blaabjerg F (2006) A new single-phase PLL structure based on second order generalized integrator. In: 2006 37th IEEE Power Electronics Specialists Conference, pp 1–6, https://doi.org/10.1109/PESC.2006.1711988

[5] Cui H, Li F, Tomsovic K (2021) Hybrid symbolic-numeric framework for power system modeling and analysis. IEEE Transactions on Power Systems 36(2):1373–1384. https://doi.org/10.1109/TPWRS.2020.3017019

[6] Dall'Anese E, Dhople SV, Johnson BB, et al (2014) Decentralized optimal dispatch of photovoltaic inverters in residential distribution systems. IEEE Transactions on Energy Conversion 29(4):957–967. https://doi.org/10.1109/TEC.2014.2357997

[7] Dugan RC, McDermott TE (2011) An open source platform for collaborating on smart grid research. In: 2011 IEEE Power and Energy Society General Meeting, pp 1–7, https://doi.org/10.1109/PES.2011.6039829

[8] Dugan RC, Montenegro D (2020) The Open Distribution System Simulator (OpenDSS) Reference Guide. Electric Power Research Institute, URL https://sourceforge.net/p/electricdss/code/HEAD/tree/trunk/Distrib/Doc/OpenDSSManual.pdf, accessed: 2026-06-01

[9] Farivar M, Neal R, Clarke C, et al (2012) Optimal inverter VAR control in distribution systems with high PV penetration. In: 2012 IEEE Power and Energy Society General Meeting, pp 1–7, https://doi.org/10.1109/PESGM.2012.6345736

[10] Hao H, Sanandaji BM, Poolla K, et al (2015) Aggregate flexibility of thermostatically controlled loads. IEEE Transactions on Power Systems 30(1):189–198. https://doi.org/10.1109/TPWRS.2014.2328865

[11] Healy W, Fanney AH, Dougherty B, et al (2017) Net Zero Energy Residential Test Facility Instrumented Data; Year 2. https://pages.nist.gov/netzero/data.html, https://doi.org/10.18434/T46W2X, minute-resolution data for February 1, 2015–January 31, 2016; accessed: 2026-05-09

[12] Hua Y, Xie Q, Hui H, et al (2021) Collaborative voltage regulation by increasing/decreasing the operating power of aggregated air conditioners considering participation priority. Electric Power Systems Research 199:107420. https://doi.org/10.1016/j.epsr.2021.107420

[13] Hua Y, Xie Q, Hui H, et al (2023) Use of inverter-based air conditioners to provide voltage regulation services in unbalanced distribution networks. IEEE Transactions on Power Delivery 38(3):1569–1579. https://doi.org/10.1109/TPWRD.2022.3218658

[14] Hui H, Ding Y, Zheng M (2019) Equivalent modeling of inverter air conditioners for providing frequency regulation service. IEEE Transactions on Industrial Electronics 66(2):1413–1423. https://doi.org/10.1109/TIE.2018.2831192

[15] Hui H, Ding Y, Chen T, et al (2021) Dynamic and stability analysis of the power system with the control loop of inverter air conditioners. IEEE Transactions on Industrial Electronics 68(3):2725–2736. https://doi.org/10.1109/TIE.2020.2975465

[16] IEEE PES Distribution System Analysis Subcommittee (2026) IEEE PES test feeders. https://cmte.ieee.org/pes-testfeeders/, accessed: 2026-05-06

[17] Kersting WH (2001) Radial distribution test feeders. In: 2001 IEEE Power Engineering Society Winter Meeting, pp 908–912, https://doi.org/10.1109/PESW.2001.916993

[18] Malhame R, Chong CY (1985) Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. IEEE Transactions on Automatic Control 30(9):854–860. https://doi.org/10.1109/TAC.1985.1104071

[19] Mathieu JL, Koch S, Callaway DS (2013) State estimation and control of electric loads to manage real-time energy imbalance. IEEE Transactions on Power Systems 28(1):430–440. https://doi.org/10.1109/TPWRS.2012.2204074

[20] National Institute of Standards and Technology (2026) Learning to use the Net-Zero Energy Residential Test Facility data. https://pages.nist.gov/netzero/learn.html, accessed: 2026-05-09

[21] Perfumo C, Kofman E, Braslavsky JH, et al (2012) Load management: Model-based control of aggregate power for populations of thermostatically controlled loads. Energy Conversion and Management 55:36–48. https://doi.org/10.1016/j.enconman.2011.10.019

[22] Ross SC, Vuylsteke G, Mathieu JL (2019) Effects of load-based frequency regulation on distribution network operation. IEEE Transactions on Power Systems 34(2):1569–1578. https://doi.org/10.1109/TPWRS.2018.2879464, URL https://arxiv.org/abs/1807.03225

[23] Turitsyn K, Sulc P, Backhaus S, et al (2011) Options for control of reactive power by distributed photovoltaic generators. Proceedings of the IEEE 99(6):1063–1073. https://doi.org/10.1109/JPROC.2011.2116750

[24] Venkatraman R, Khaitan SK, Ajjarapu V (2019) Dynamic co-simulation methods for combined transmission-distribution system with integration time step impact on convergence. IEEE Transactions on Power Systems 34(2):1171–1181. https://doi.org/10.1109/TPWRS.2018.2874807

[25] Zhao N, Gorbachev S, Yue D, et al (2022) The effect of communication delays on the frequency stability of power systems integrated with inverter air conditioners. Sustainable Energy, Grids and Networks 32:100920. https://doi.org/10.1016/j.segan.2022.100920

[26] Zhu GX, Bao YQ, Yu QQ (2024) A control strategy for air-conditioning loads participating in frequency regulation based on model predictive control. Sustainable Energy, Grids and Networks 38:101369. https://doi.org/10.1016/j.segan.2024.101369

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM1.pdf
- ESM2.zip