

A three-layer hierarchical dynamic operating envelope framework for residential distributed energy resource integration in unbalanced low-voltage distribution networks

Omar Alrumayh

o.alrumayh@qu.edu.sa

Qassim University

Article

Keywords:

Posted Date: October 23rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11212759/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

A three-layer hierarchical dynamic operating envelope framework for residential distributed energy resource integration in unbalanced low-voltage distribution networks

Omar Alrumayh^{1,*}

¹Department of Electrical Engineering, College of Engineering, Qassim University, Buraydah, 52571, Saudi Arabia

*o.alrumayh@qu.edu.sa

ABSTRACT

Increasing penetration of residential distributed energy resources (DERs) drives low-voltage feeders beyond their voltage and thermal limits. Existing hierarchical schemes that deliver dynamic operating envelopes to households typically bound import only, neglect the energy deferred by curtailment, and issue bounds that households cannot be guaranteed to meet. This paper proposes a three-layer framework. Each home computes a cost-optimal and an energy-minimal schedule; an aggregator separates enforceable import- and export-reduction capacity from voluntary recovery headroom and issues hard import and export limits with energy-bounded payback, each anchored on the energy-minimal dispatch and therefore provably feasible; and the network operator sizes the relief with an unbalanced AC optimal power flow and verifies it in OpenDSS, tightening envelopes only in violated intervals. On the IEEE European Low-Voltage Test Feeder under a commercial export-paying tariff, the binding event is a coincident reverse-power peak that import-only envelopes cannot address. The proposed framework is the only scheme tested that sustains 100% DER penetration within both voltage and thermal limits, compared with 25% and 50% for the benchmark schemes. It reduces peak loading by 40% and network losses by 33% relative to uncoordinated operation, without relaxing comfort constraints or leaving any household worse off after settlement.

Nomenclature

Sets, parameters and decision variables used in the paper.

Symbol	Description
<i>Sets, indices and labels</i>	
$\mathcal{T}, t; \mathcal{H}, h$	Time steps; households
$\mathcal{N}, k; \mathcal{P}, p; \mathcal{L}, \ell$	Buses (substation $k = s$); phases $\{a, b, c\}$; lines
$\mathcal{H}_{k,p}$	Households connected at node (k, p)
$\mathcal{F}, n$	Feeders
$\mathcal{A}, a; \mathcal{D}$	Household assets; sheddable or shiftable device actions
$s \in \{\text{BESS}, \text{EV}\}$	Storage devices
$x \in \{\text{HP}, \text{fr}\}$	Thermostatic loads (heat pump, refrigerator)
$\mathcal{W}, w$	Wet appliances $\{\text{W/D}, \text{DW}\}$
$\mathcal{S}_h$	Plug-in session of vehicle h (wraps midnight)
$\mathcal{C}, c$	Softened network limits in the AC-OPF
Π, v	DER penetration levels; penetration level [%]
i	Iteration of the AC-verified tightening loop
CO, EM, RD	Cost-optimal, energy-minimal and re-dispatch HEMS solves
<i>Parameters</i>	
$T, \Delta t$	Number of steps; step length [h]
$\lambda_t^{\text{imp}}, \lambda_t^{\text{exp}}$	Import / export prices [£/kWh]
$\lambda_{h,t}^{\text{inc}}, \lambda_{h,t}^{\text{fee}}$	Case-2 compliance rate; Case-3 service fee [£/kWh]
$\bar{P}^{\text{imp}}, \bar{P}^{\text{exp}}, M$	Connection import / export limits; big- M constant [kW]
$P_{\text{PV},h,t}^{\text{av}}$	Available PV generation [kW]
$\bar{E}_h^s, \underline{\sigma}^s, \sigma^{\text{end}}$	Storage capacity [kWh]; SoC floor fraction; BESS initial and terminal SoC fraction
$\bar{P}_h^{\text{ch},s}, \bar{P}_h^{\text{dis},s}$	Storage charge / discharge ratings [kW]

Symbol	Description
$\eta_h^{\text{ch},s}, \eta_h^{\text{dis},s}, \eta_{s,h}^{\pm}$	One-way charge / discharge efficiencies; efficiencies in the SoC balance
$t_h^{\text{arr}}, t_h^{\text{dep}}; E_h^{\text{EV},0}, E_h^{\text{dep}}$	EV arrival / departure steps; arrival / required departure energy [kWh]
$R_h^{\text{EWH}}, C_h^{\text{EWH}}, Q_h^{\text{EWH}}$	EWH tank thermal resistance [$^{\circ}\text{C}/\text{kW}$] and capacitance [$\text{kWh}/^{\circ}\text{C}$]; element rating [kW]
$M_h^{\text{EWH}}, m_{h,t}$	Tank capacity; hot-water draw in step t [L]
$R_h^x, C_h^x, \text{COP}_h^x, \alpha_h^x$	Thermal resistance [$^{\circ}\text{C}/\text{kW}$], capacitance [$\text{kWh}/^{\circ}\text{C}$], COP, retention factor
$T_{h,t}^a, T_t^{\text{amb}}, T_{h,t}^{\text{kit}}$	Tank-room, outdoor and kitchen air temperatures [$^{\circ}\text{C}$]
$\underline{T}_h^{\bullet}, \bar{T}_h^{\bullet}$	Lower / upper comfort limits of thermal device $\bullet$ [$^{\circ}\text{C}$]
$\bar{P}_h^x, \bar{P}_h^w$	Heat-pump / refrigerator and wet-appliance electrical ratings [kW]
$D_h^w, [t_h^{\text{e},w}, t_h^{\text{l},w}], \chi_h^w$	Wet-appliance cycle length; permitted window; Bernoulli requirement indicator
$w_{\text{trk}}, w_{\text{env}}$	Import-target tracking weight; envelope-slack penalty [$\text{£}/\text{kWh}$]
$w_{\text{need}}, w_{\text{extra}}, w_{\text{pk}}$	Tier-one and tier-two payback priorities; feeder-peak penalty
$\gamma, \eta^{\text{rec}}$	Fairness strength; recoverable fraction of curtailed energy
E_h^{shift}	Total shiftable energy of home h [kWh]
$\bar{S}_{\text{Dx}}^{\text{loss}}, \bar{S}_{\text{Dx}}$	Loss-adjusted and nameplate substation ratings [kVA]
$\bar{P}_n^{\text{base}}$	Uncoordinated peak of feeder n [kW]
PF	Assumed constant power factor
$w_F, w_L, w_S; w_F^{(2)}$	AC-OPF weights on flexibility, losses and violation slack; Case-2 weight on the signed request
$\mathbf{Y}, \mathbf{Y}_{kk'}$	Three-phase bus admittance matrix; 3×3 off-diagonal block for line $k-k'$
$\underline{V}, \bar{V}, V^{\text{ref}}, \theta_p$	Voltage limits; substation reference voltage; phase reference angle [p.u.]
$\bar{I}_{\ell}, \kappa$	Ampacity of line ℓ [p.u.]; fraction of ampacity the request is sized against
ψ^{flex}	Reactive-to-active ratio of the flexibility request
β, K_{max}	Geometric tightening factor; maximum number of dispatch–replay passes
ε	Optional numerical margin on the import cap (zero in this study)
μ^{fl}, ξ_h	Case-2 deployment floor fraction; per-home share of the node request
<i>Variables and derived quantities</i>	
$P_{h,t}^{\text{imp}}, P_{h,t}^{\text{exp}}, P_{h,t}^{\text{net}}$	Grid import, export and net power $P_{h,t}^{\text{imp}} - P_{h,t}^{\text{exp}}$ of home h [kW]
$X_{h,t}$	Import–export complementarity binary
$P_{a,h,t}; P_{\text{PV},h,t}, P_{\text{curt},h,t}$	Power of asset a (positive when drawing); utilised / curtailed PV [kW]
$\mathbf{u}_h$	Decision variables of the HEMS of home h
$\mathbf{P}_h^{\text{CO}}, \mathbf{P}_h^{\text{EM}}$	Cost-optimal / energy-minimal schedules of home h
$P_{h,t}^{\bullet, \text{CO}}, P_{h,t}^{\bullet, \text{EM}}$	Quantity $\bullet$ in the CO / EM schedule; $P_{h,t}^{\text{imp}, \text{EM}}$ and $P_{h,t}^{\text{exp}, \text{EM}}$ are the feasibility floors
$E_{h,t}^s; P_{h,t}^{\text{ch},s}, P_{h,t}^{\text{dis},s}$	Storage SoC [kWh]; charge / discharge power [kW]
$x_{h,t}^s, u_{h,t}^s$	Storage mode binary; availability flag
$T_{h,t}^{\text{EWH}}, T_{h,t}^{\text{EWH}}, P_{h,t}^{\text{EWH}}$	Tank temperature [$^{\circ}\text{C}$]; element on–off state; EWH electrical power [kW]
$T_{h,t}^x, \phi_{h,t}^{x,+}, \phi_{h,t}^{x,-}$	Indoor or cabinet temperature [$^{\circ}\text{C}$]; heating / cooling duty fractions
$x_{h,t}^w, y_{h,t}^w$	Wet-appliance running / start indicators
$s_{h,t}^{\text{DOE}}$	Import-envelope slack (fixed at zero in the proposed scheme) [kW]
$P_{h,t}^{\text{imp}, \text{tgt}}$	Aggregator per-home import target [kW]
$f_{h,t}^{\text{imp}}, f_{h,t}^{\text{exp}}, f_{h,t}^{\text{rec}}$	Import- and export-reduction capacity; voluntary recovery headroom [kW]
$\Delta P_{h,t}^{\text{dis},s}, \Delta P_{h,t}^{\text{ch},s}, \Delta P_{h,t}^{\text{EWH}, \uparrow}, \Delta P_{h,t}^{\text{HP}, \uparrow}$	Spare storage discharge / charge power; spare EWH / HP power [kW]
$F_{k,p,t}^{\text{imp}}, F_{k,p,t}^{\text{exp}}$	Node import- / export-reduction bundles [kW]
$\delta_{h,t}^{\text{imp}}, \delta_{h,t}^{\text{exp}}$	Allocated import / export cuts [kW]
$r_{h,t}, r_{h,t}^{(1)}, r_{h,t}^{(2)}$	Payback recovery granted, total / tier one / tier two [kW]
$\pi_n; E_h^{\text{cut}}, P_{h,t}^{\text{cut}}$	Feeder peak [kW]; curtailed energy [kWh]; post-curtailement net power [kW]
$\omega_h, \vartheta_h, \bar{\vartheta}$	Fairness weight; curtailment burden ratio; fleet-mean burden ratio
$\bar{P}_{h,t}^{\text{DOE}}, \underline{P}_{h,t}^{\text{DOE}}$	Import cap; export cap ($-\underline{P}_{h,t}^{\text{DOE}}$) of home h [kW]
$\bar{P}_{h,t}^{(i)}, \underline{P}_{h,t}^{(i)}$	Magnitude of a tightened cap at iteration i ; its feasibility floor [kW]
B_h, B_h^{CO}	Settled and uncoordinated daily bills of home h [£]
$p_{k,p,t}^{\text{flex}}, Q_{k,p,t}^{\text{flex}}, \zeta_{k,p,t}$	AC-OPF flexibility request; its reactive part; epigraph of $ p_{k,p,t}^{\text{flex}} $
$p_{k,p,t}^d, Q_{k,p,t}^d, F_{k,p,t}^{\text{exp}}$	Aggregated demand and export at node (k, p)
$P_{k,p,t}^s, Q_{k,p,t}^s, S_{k,p,t}^s$	Substation injections (non-zero only at $k = s$)
$V_{k,p,t}, \mathbf{V}_t, \mathbf{I}_t$	Nodal phase voltage ($V^{\text{re}} + jV^{\text{im}}$); voltage and injected-current vectors [p.u.]
$I_{\ell,p,t}, \mathbf{I}_{\ell,t}$	Branch current of line ℓ [p.u.]
$P_{\text{loss},t}, s_c$	Network active-power loss [kW]; violation slack
$F_{k,p,t}^{(2)}, P_{h,t}^{\text{gross}, \text{EM}}$	Case-2 node bundle; gross controllable load in the EM schedule [kW]

1 Introduction

The decarbonisation of household energy is reshaping low-voltage (LV) distribution networks. Rooftop photovoltaic (PV) systems, behind-the-meter batteries, electric vehicles (EVs) and electrified heating are being connected at a scale for which distribution networks were not designed. In the United Kingdom, Australia and Germany, national planning scenarios and expansion targets foresee rapid growth of rooftop PV, electric vehicles and heat pumps on residential connections over the coming decade [1–3]. This shift turns passive consumers into *prosumers* that both draw from and inject into the feeder, a trend reinforced by the move from feed-in tariffs toward self-consumption schemes [3, 4]. Although these resources can shift demand, store surplus generation and support the grid, their uncoordinated operation can adversely affect the network. When midday PV export coincides with light local load, or when least-cost charging synchronises in the cheapest hours [5], bus voltages exceed the EN 50160 ceiling (1.10p.u.) [6] and cables and transformers are driven beyond their thermal ratings, accelerating asset ageing and degrading reliability [7, 8].

1.1 Background and motivation

The home energy management system (HEMS) schedules household devices to minimise the prosumer’s bill while preserving comfort. However, a HEMS is *network-unaware*: it optimises a single household without visibility of the feeder state [9]. Two difficulties follow. First, individually rational schedules synchronise onto the hours that the tariff rewards (the cheapest import hours and, under export-paying tariffs, the highest-priced export hours) and combine into a collectively infeasible operating point [10]. Second, a fully centralised controller that embeds the AC optimal power flow (OPF) in a large mixed-integer nonlinear program scales poorly with the number of dwellings and requires intrusive device-level metering. Network-side alternatives based on reconfiguration or dynamic tariffs [11], or on transactive mechanisms [12], avoid device-level scheduling but place the optimisation burden on the operator or a market and offer no explicit guarantee at the connection point. Reconciling the cost and comfort objectives of the HEMS with the voltage and thermal objectives of the distribution system operator (DSO) therefore requires an explicit coordination layer [13], and device-level modelling is needed to reveal the shiftable energy on which such a layer can draw [14].

Network reinforcement is capital-intensive and frequently leaves assets oversized for most of their service life [15]. A less costly alternative is to coordinate the flexibility already available behind the meter, whose definition and provision mechanisms vary widely in the literature [16]. In the approach adopted here, the operator communicates per-household *dynamic operating-envelope* bounds to each HEMS, retaining the scalability and privacy of distributed control while providing sufficient network awareness. The open question is how to compute envelopes that are tight enough not to restrict the hosting of distributed energy resources (DERs) unnecessarily and reliable enough that universal compliance guarantees a feasible feeder, without penalising the households that provide the flexibility. This *hosting-capacity allocation* problem is the subject of the present paper.

1.2 Literature review and related work

Hosting capacity (HC) quantifies the amount of DER a feeder can accommodate before a voltage or thermal limit is reached. Early static methods assume worst-case coincident generation and are therefore conservative [17]. Probabilistic time-series assessment [8] and probabilistic load flow [18] replaced this worst case with diversity-aware headroom, while studies of PV reactive-power control [7] and voltage-stability limits [19] treated headroom as a function of location and time; a survey is given in [20]. Sensitivity-based screening built on the DistFlow model [21] and reactive-power sensitivities [22] provides fast analytic estimates, and receding-horizon methods size distributed storage against the resulting headroom [23]. In contrast, the present work defines hosting capacity operationally as the highest penetration level sustained without violation, as determined by an AC-verified envelope-tightening loop.

Dynamic operating envelopes (DOEs), also termed customer operating envelopes [24], map feeder-level HC onto per-customer import and export limits. Sensitivity-based DOEs have been shown to maintain voltage compliance [25], OPF-based envelopes with import and export limits have been used to facilitate residential DER services [26], and residential DERs have been coordinated through an affinely adjustable robust ADMM scheme that remains feasible under forecast deviations [27]. The sensitivity coefficients required on unbalanced radial feeders were derived in [28], and the underlying OPF was decentralised using learned local controllers in [29]. Recent reviews classify DOE computation into optimisation-based, iterative power-flow and data-driven methods, and identify unbalanced modelling, forecast accuracy and fairness as open problems [30, 31], while Australian programmes are standardising the envelope definition and its customer interface [32]. Lankeshwara et al. [33] experimentally validated a DOE-enabled demand-response scheme in which the network operator issues envelopes and an aggregator keeps household loads within them; this is the closest counterpart to the present architecture, and it likewise bounds the customer from the network side only. In general, envelopes are computed on the network side and passed down, whereas whether the household can re-dispatch within them, and at what cost, is treated separately or delegated to a price signal. The present paper provides a complementary customer-side guarantee: each cap is anchored on the home’s maximum-self-consumption dispatch, so a feasible re-dispatch provably exists, and this anchor is coupled with an

energy-bounded payback and an ex-post settlement that guarantees that no participating household is worse off.

Hierarchical control organises DER coordination into layers with separated time scales and information scope. Early schemes are two-level, pairing home-level scheduling with an aggregator or DSO layer, for example through transactive EV charging [34] or dynamic pricing for demand response [35], but they do not embed network limits; distributed OPF methods suited to the aggregator layer are surveyed in [36]. Other two-level work optimises resources without an explicit network layer [37]. More recent three-level architectures span the supply, retailer and demand sides [38], a DSO, a microgrid and its consumers [39], an energy hub, aggregator and end users [40], and grid-aware hierarchical coordination on realistic distribution feeders [41]. Most closely related to the present work, Mak and Choi [42] and Hussain et al. [43] placed a HEMS, an aggregator and a volt-var or OPF layer in series, reporting lower losses and prosumer bills, and decentralised, centralised and hierarchical coordination of residential DERs has been compared on an unbalanced LV network [44]. Market-oriented three-layer frameworks have also been developed, in which an aggregator prices household flexibility [45] or smart homes trade energy subject to network constraints [46]. These studies show that an intermediate layer can reconcile household and network objectives, but they procure this reconciliation through prices or trading positions, so network limits are respected in expectation rather than bounded at the connection point.

Two limitations recur across this body of work. First, although a DOE by definition specifies both an import limit and an export limit at the connection point [24, 32], hierarchical HEMS-coordination frameworks almost invariably mobilise import-reduction flexibility only; in effect, they enforce the import limit and leave the export limit unconstrained, so they neither bound PV-driven reverse power flow nor coordinate import and export during the coincident peak. Second, where flexibility is priced, compliance relies on an incentive that guarantees neither feeder feasibility nor full compensation of the participating households [47–49]. Recent reviews confirm that envelope computation is still posed as a network-side problem, with household deliverability and cost allocation treated separately [30, 31].

In the proposed framework, the DSO only sizes and verifies the relief request, while the device-level dispatch remains in the home, so no device telemetry leaves the dwelling. The scheme is a sequential hierarchical coordination rather than a nested bi-level program: each layer solves its own optimisation, and no layer optimises over another's reaction function. For the HEMS, tractable device models are adopted for EV charging [50], battery storage [51], water heating [52], space conditioning [53], and deferrable or cycling loads [54, 55], within the mixed-integer linear programming (MILP) approach standard for day-ahead scheduling [56].

1.3 Research gap and contribution

Three gaps remain at the interface between the envelope and the household that must respect it. (i) Although DOEs comprise both import and export limits, hierarchical HEMS-coordination frameworks typically mobilise import-side flexibility only, so the export limit of the envelope is rarely used as a control variable. (ii) Envelopes derived from instantaneous flexibility signals overlook the fact that curtailment shifts *energy*; unless deferred energy is returned in a network-safe window, it rebounds and may create a new peak. (iii) Envelope construction and household re-dispatch are coupled only through a price incentive, so there is no guarantee that a home can meet its bounds without relaxation, nor that participating households are fully compensated.

This paper addresses these gaps with a three-layer scheme based on feasibility-anchored, energy-bounded envelopes acting on both import and export. The main contributions are as follows.

1. An integrated three-layer coordination framework (HEMS, envelope aggregator, and DSO AC-OPF with AC-verified tightening), evaluated on the IEEE European LV Test Feeder over a 0–100% DER penetration sweep. It is the only scheme tested that sustains 100% penetration within both the EN 50160 voltage band and conductor ampacity, compared with hosting ceilings of 25% for the decentralised baseline and 50% for a conventional import-only hierarchical scheme.
2. A state-aware decomposition of each home's flexibility into what the aggregator may *enforce* (import- and export-reduction capacity) and what the home only *offers* (voluntary recovery headroom), evaluated along the home's scheduled state trajectory and bounded by its energy-minimal schedule.
3. Hard operating envelopes, with both import and export limits enforced, that carry *provable feasibility floors*: since each cap dominates the home's energy-minimal dispatch, the re-dispatch is always feasible with zero envelope slack, making compliance a guarantee rather than an incentivised target.
4. An AC-verified tightening loop that shrinks envelopes only at violated time steps and terminates either violation-free or at the feeder's physical hosting-capacity boundary.
5. The identification of export-limit enforcement, rather than flexibility volume, as the source of the hosting gain, together with a transparent cost accounting. At full penetration the framework requests 86% of the conventional scheme's

import-side flexibility, additionally requests 143 kWh/day of export reduction, reduces peak loading by 38% relative to that scheme and network losses by 33% relative to the uncoordinated case, with zero PV curtailment and no comfort violation. The resulting net bill premium (£0.74 per home per day), which stems from forgone export revenue, is reimbursed to households through an ex-post settlement.

1.4 Organisation of the paper

Section 2 presents the network model, the models of the three control layers, the coordination architecture and algorithm, and the controller scenarios. Section 3 presents the simulation results, the ablation study and the limitations. Section 4 concludes the paper.

2 Methodology

2.1 Overview

This section first presents the network model (Section 2.2) and the models of the three control layers: Layer 1, the home-level HEMS (Section 2.3); Layer 2, the aggregator envelopes (Section 2.4); and Layer 3, the DSO-level OPF (Section 2.5), followed by the hosting-capacity metric (Section 2.6). It then describes the coordination architecture (Sections 2.7–2.11), the operational sequence and its convergence (Section 2.12), and the controller scenarios (Section 2.13).

Each participating dwelling is assumed to have a smart meter, a two-way communication link to the aggregator and a HEMS. DER owners hold a common bundle comprising a rooftop PV array, a battery energy storage system (BESS) and an EV with vehicle-to-grid (V2G) capability. In addition, every household operates flexible loads: an electric water heater (EWH), a reversible air-source heat pump (HP), two uninterruptible wet appliances (a washer/dryer, W/D, and a dishwasher, DW) and a refrigerator; the remaining demand is treated as inflexible base load. Because the second group is present at every penetration level, flexibility is available even at 0% DER. Decisions are taken over a day-ahead horizon of $T = 48$ half-hourly steps ($\Delta t = 0.5$ h), indexed by $t \in \mathcal{T} = \{1, \dots, T\}$; households are indexed by $h \in \mathcal{H}$, network nodes by the bus–phase pair (k, p) , and $\mathcal{H}_{k,p}$ denotes the homes connected at node (k, p) . $P_{h,t}^{\text{imp}}$ and $P_{h,t}^{\text{exp}}$ denote the grid import and export of home h , and all constraints hold for all $t \in \mathcal{T}$ and $h \in \mathcal{H}$ unless stated otherwise.

2.2 Network model and power-flow formulation

The feeder is modelled as an unbalanced three-phase LV network with bus set $\mathcal{N}$, phase set $\mathcal{P} = \{a, b, c\}$ and line set $\mathcal{L}$. Network feasibility requires *voltage limits* $\underline{V} \leq |V_{k,p,t}| \leq \bar{V}$, with $\underline{V} = 0.90$ p.u. and $\bar{V} = 1.10$ p.u. (EN 50160 LV band, $\pm 10\%$); and *thermal limits* on line current, $|I_{\ell,p,t}| \leq \bar{I}_{\ell}$, and on transformer apparent power, $S_{\text{tx},t} \leq \bar{S}_{\text{tx}}$. Power flows are solved natively in OpenDSS with a Newton–Raphson solver and a constant-power (PQ) load representation; the coupling between the optimisation layer and OpenDSS follows [57].

2.3 Layer 1: home energy management

Each household h solves a day-ahead mixed-integer linear program (MILP) over the 48-step horizon, minimising the prosumer’s electricity cost subject to the operating-envelope bounds broadcast by the aggregator and the comfort limits of every device. Let $\mathbf{u}_h$ collect the home’s decision variables (device powers, storage charge and discharge, and the binary mode indicators).

Layer 1 is solved in three stages that share the identical device constraint set (5)–(21) and differ only in the objective: a cost-optimal reference solve (CO), an energy-minimal companion solve (EM) that supplies the feasibility floors of Section 2.4.3, and a re-dispatch solve (RD) executed after the aggregator broadcasts the hard envelopes.

2.3.1 Cost-optimal reference schedule (CO)

The reference solve minimises the household’s net energy bill under the time-of-use tariff,

$$\min_{\mathbf{u}_h} \sum_{t \in \mathcal{T}} \left(\lambda_t^{\text{imp}} P_{h,t}^{\text{imp}} - \lambda_t^{\text{exp}} P_{h,t}^{\text{exp}} \right) \Delta t, \quad (1)$$

where λ_t^{imp} and λ_t^{exp} are the import and export prices. The first term is the cost of energy drawn from the grid and the second the revenue from energy exported to it. No regularising penalties are applied: the device constraints alone determine the schedule, so (1) represents purely bill-driven, uncoordinated behaviour. Its solution is the cost-optimal schedule $\mathbf{P}_h^{\text{CO}}$, which serves both as the decentralised baseline (Case 1) and as the reference against which flexibility is measured in Section 2.4.

2.3.2 Energy-minimal companion schedule (EM)

The companion solve replaces the price-weighted objective with an unpriced energy objective,

$$\min_{\mathbf{u}_h} \sum_{t \in \mathcal{T}} \left(P_{h,t}^{\text{imp}} + P_{\text{curt},h,t} \right) \Delta t, \quad (2)$$

in which $P_{\text{curt},h,t}$ is curtailed PV power. Penalising grid import and curtailment together drives the home to absorb its own generation locally rather than to import cheaply or spill surplus, so (2) yields the maximum-self-consumption dispatch. Because it is solved under the same hard device and comfort constraints, the resulting schedule $\mathbf{P}_h^{\text{EM}}$ is the deepest reduction in grid exchange the home can physically reach, and it supplies the feasibility floors $P_{h,t}^{\text{imp,EM}}$.

2.3.3 Re-dispatch under the hard envelopes (RD)

Once Layer 2 has broadcast the per-home caps, the household re-solves

$$\min_{\mathbf{u}_h} \sum_{t \in \mathcal{T}} \left[\left(\lambda_t^{\text{imp}} P_{h,t}^{\text{imp}} - \lambda_t^{\text{exp}} P_{h,t}^{\text{exp}} \right) - \lambda_t^{\text{inc}} P_{h,t}^{\text{imp}} + w^{\text{env}} s_{h,t}^{\text{DOE}} + w^{\text{trk}} \left| P_{h,t}^{\text{imp}} - P_{h,t}^{\text{imp,tgt}} \right| \right] \Delta t. \quad (3)$$

The first bracket is the energy bill of (1). The second term is a per-kWh compliance rate $\lambda_{h,t}^{\text{inc}}$ on imported energy and the third penalises the import-envelope slack $s_{h,t}^{\text{DOE}}$; both are used only by the conventional benchmark (Section 2.13.1). The proposed scheme sets the rate to zero, compensates households ex post (Section 2.4.4) and fixes the slack at zero, so its caps are hard. The final term, used only by the proposed scheme, is a tracking penalty of weight w^{trk} that pulls the realised import toward the aggregator's per-home target $P_{h,t}^{\text{imp,tgt}}$.

The tracking weight must be chosen with care. Since w^{trk} penalises each kWh of deviation, a value above the import price would make it cheaper for a household to purchase unneeded energy than to remain below its target, turning the target into an instruction to consume. The tracking weight is therefore required to satisfy

$$w^{\text{trk}} < \min_t \lambda_t^{\text{imp}}, \quad (4)$$

and $w^{\text{trk}} = 0.05$ £/kWh is adopted, against the lowest import price of 0.1517 £/kWh (Section 3.1). Under (4), the term resolves ties among schedules of equal cost without rendering wasteful consumption profitable.

2.3.4 Power balance and grid limits

The household power balance couples every device to the grid connection point. The following constraints hold $\forall t \in \mathcal{T}$ and $\forall h \in \mathcal{H}$:

$$P_{h,t}^{\text{imp}} - P_{h,t}^{\text{exp}} = \sum_{a \in \mathcal{A}} P_{a,h,t} - P_{\text{PV},h,t}, \quad (5)$$

$$0 \leq P_{h,t}^{\text{imp}} \leq \bar{P}^{\text{imp}}, \quad 0 \leq P_{h,t}^{\text{exp}} \leq \bar{P}^{\text{exp}}, \quad (6)$$

$$P_{\text{PV},h,t} = P_{\text{PV},h,t}^{\text{av}} - P_{\text{curt},h,t}, \quad 0 \leq P_{\text{curt},h,t} \leq P_{\text{PV},h,t}^{\text{av}}, \quad (7)$$

$$\underline{P}_{h,t}^{\text{DOE}} \leq P_{h,t}^{\text{net}} \leq \bar{P}_{h,t}^{\text{DOE}} + s_{h,t}^{\text{DOE}}. \quad (8)$$

Here $\mathcal{A} = \{\text{base, EV, BESS, EWH, HP, W/D, DW, fr}\}$ indexes the household assets, namely the inflexible base load, the EV, the BESS, the EWH, the HP, the W/D, the DW and the refrigerator, with $P_{a,h,t}$ positive when the asset draws power and negative when it injects. Equation (5) states the power balance at the connection point, and (7) splits the available PV generation $P_{\text{PV},h,t}^{\text{av}}$ into the utilised part and the curtailed part $P_{\text{curt},h,t}$. Constraint (6) bounds import and export by $\bar{P}^{\text{imp}} = \bar{P}^{\text{exp}} = 100$ kW, equal to the big- M constant below, so these bounds never bind and the connection is limited only by the envelope, and (8) enforces the operating envelope on the net power $P_{h,t}^{\text{net}} = P_{h,t}^{\text{imp}} - P_{h,t}^{\text{exp}}$ between the export cap $\underline{P}_{h,t}^{\text{DOE}}$ and the import cap $\bar{P}_{h,t}^{\text{DOE}}$ (constructed in Section 2.4.3). The two sides are enforced differently. The *export* cap admits no slack, which is possible because the floor $P_{h,t}^{\text{exp,EM}}$ of (28b) guarantees a feasible point beneath it. The *import* cap carries a slack $s_{h,t}^{\text{DOE}} \geq 0$ in the general formulation; the proposed scheme fixes it at zero, which the floor of (28a) makes possible, whereas the conventional benchmark keeps it and penalises it at w^{env} . A big- M complementarity, $P_{h,t}^{\text{imp}} \leq MX_{h,t}$ and $P_{h,t}^{\text{exp}} \leq M(1 - X_{h,t})$ with binary $X_{h,t}$, prevents simultaneous import and export.

2.3.5 Battery and electric-vehicle storage

The BESS and the EV are modelled as one class of store, indexed by $s \in \{\text{BESS}, \text{EV}\}$, with state of charge (SoC) $E_{h,t}^s$ and charge and discharge powers $P_{h,t}^{\text{ch},s}$ and $P_{h,t}^{\text{dis},s}$:

$$E_{h,t}^s = E_{h,t-}^s + \left(\eta_{s,h}^+ P_{h,t}^{\text{ch},s} - \frac{P_{h,t}^{\text{dis},s}}{\eta_{s,h}^-} \right) \Delta t, \quad (9)$$

$$0 \leq P_{h,t}^{\text{ch},s} \leq \bar{P}_h^{\text{ch},s} u_{h,t}^s x_{h,t}^s, \quad 0 \leq P_{h,t}^{\text{dis},s} \leq \bar{P}_h^{\text{dis},s} u_{h,t}^s (1 - x_{h,t}^s), \quad (10)$$

$$\underline{\sigma}^s \bar{E}_h^s \leq E_{h,t}^s \leq \bar{E}_h^s, \quad (11)$$

where t^- is the step preceding t , $u_{h,t}^s \in \{0, 1\}$ marks availability, the binary $x_{h,t}^s$ forbids simultaneous charge and discharge, and the net power $P_{s,h,t} = P_{h,t}^{\text{ch},s} - P_{h,t}^{\text{dis},s}$ enters (5).

For the battery, $u_{h,t}^{\text{BESS}} \equiv 1$, $t^- = t - 1$, $\eta^+ = \eta_h^{\text{ch},\text{BESS}} = 0.86$ and $\eta^- = \eta_h^{\text{dis},\text{BESS}} = 0.85$ (round trip 0.73), and $\underline{\sigma}^{\text{BESS}} = 0.10$. Each home draws its usable capacity $\bar{E}_h^{\text{BESS}}$ from a domestic pool (5.0–13.5 kWh, power rating auto-scaled to 2.5–5.0 kW) anchored on the Tesla Powerwall 2 [58], and the daily cycle is closed by

$$E_{h,T}^{\text{BESS}} \geq E_{h,0}^{\text{BESS}} = \sigma^{\text{end}} \bar{E}_h^{\text{BESS}}, \quad (12)$$

with $\sigma^{\text{end}} = 0.50$. The condition is an inequality so that a home is never forced to discard stored energy; it is active in every solved instance, so the battery is energy-neutral over the day.

The EV is controllable only during its plug-in session $\mathcal{S}_h \subset \mathcal{T}$, the ordered set of steps from its arrival t_h^{arr} to its departure t_h^{dep} , which wraps through midnight because the vehicle arrives in the evening and leaves the following morning. Accordingly, $u_{h,t}^{\text{EV}} = 1$ only on $\mathcal{S}_h$, t^- is the preceding step within $\mathcal{S}_h$, and the EV variables are fixed to zero outside it. The SoC is expressed on a grid-equivalent basis, with the round-trip efficiency applied on charging, $\eta^+ = \eta_h^{\text{ch},\text{EV}} \eta_h^{\text{dis},\text{EV}}$ and $\eta^- = 1$, and $\underline{\sigma}^{\text{EV}} = 0.20$. The session boundaries are

$$E_{h,t_h^{\text{arr}}}^{\text{EV}} = E_h^{\text{EV},0}, \quad E_{h,t_h^{\text{dep}}}^{\text{EV}} = \underline{E}_h^{\text{dep}}, \quad (13)$$

where the departure requirement is 0.80 of capacity, which the 13.5 h session can always reach at 3.6 kW; the equality prevents EV energy from being carried across the horizon boundary. The charger is a 3.6 kW single-phase unit, the V2G discharge rating is 3.6 kW in 40 homes and 2.8 kW in 15, battery capacities range from 19 to 52 kWh, and the one-way efficiencies are $\eta_h^{\text{ch},\text{EV}} \in \{0.89, 0.91, 0.94\}$ and $\eta_h^{\text{dis},\text{EV}} \in \{0.91, 0.92, 0.95\}$ by vehicle type. All vehicles share a common plug-in session from 18:30 to 08:00, representative of evening arrival and morning departure in commuter travel [59]; holding the session common, like the initial states of charge (Section 3.2), prevents storage-driven results from being attributed to a favourable draw. Because the session spans midnight, the 02:00–05:00 low-price window lies within the plug-in period of every vehicle, which explains why least-cost scheduling concentrates charging there; the vehicles are away during the 16:00–19:00 export band.

2.3.6 Electric water heater

The EWH is a thermal store, modelled with the first-order tank model of [60–62], in which the water temperature relaxes toward the air of the space housing the tank, the relaxation being strengthened whenever hot water is drawn off and replaced:

$$T_{h,t}^{\text{EWH}} = T_{h,t}^{\text{a}} + R_h^{\text{EWH}} Q_h^{\text{EWH}} z_{h,t}^{\text{EWH}} - \frac{M_h^{\text{EWH}} - m_{h,t}}{M_h^{\text{EWH}}} (T_{h,t}^{\text{a}} - T_{h,t-1}^{\text{EWH}}) e^{-\Delta t / (R_h^{\text{EWH}} C_h^{\text{EWH}})}, \quad (14)$$

$$\underline{T}_h^{\text{EWH}} \leq T_{h,t}^{\text{EWH}} \leq \bar{T}_h^{\text{EWH}}, \quad T_{h,T}^{\text{EWH}} \geq T_{h,0}^{\text{EWH}}, \quad (15)$$

$$P_{h,t}^{\text{EWH}} = Q_h^{\text{EWH}} z_{h,t}^{\text{EWH}}, \quad z_{h,t}^{\text{EWH}} \in \{0, 1\}. \quad (16)$$

In (14), Q_h^{EWH} is the rating of the resistive element (3.0–3.6 kW), $z_{h,t}^{\text{EWH}}$ its on–off state, R_h^{EWH} the thermal resistance and C_h^{EWH} the thermal capacitance of the tank. The second term is the temperature rise contributed by the element in the step ($R_h^{\text{EWH}} = 2.03\text{--}3.04^\circ\text{C/kW}$, inversely proportional to the tank volume, so that one half-hour step at full power raises the tank by 6–11°C). The third term is the standing-loss relaxation of the previous temperature toward the ambient $T_{h,t}^{\text{a}}$, with time constant $R_h^{\text{EWH}} C_h^{\text{EWH}} = 98$ h, scaled by the fraction of the tank not displaced by the hot-water draw $m_{h,t}$; $T_{h,0}^{\text{EWH}}$ is the initial tank temperature. Replacement water is assumed to enter at ambient rather than mains temperature, which understates the reheat energy and makes the EWH a conservatively *smaller* flexibility resource. The comfort band (15) is set per home (lower limits of 46–51°C and upper limits of 58–62°C), and the tank must end the day at least at its initial temperature. The stored heat allows the HEMS to pre-heat when energy is cheap or when the network requires relief.

2.3.7 Controllable uninterruptible demand (wet appliances)

The household operates two wet appliances, indexed by $w \in \mathcal{W} = \{W/D, DW\}$ for the washer/dryer and the dishwasher. Each is a deferrable but uninterruptible load: once started, its cycle runs without interruption for D_h^w steps and must complete inside the user's permitted window $[t_h^{e,w}, t_h^{l,w}]$,

$$x_{h,t}^w = \sum_{\tau=t-D_h^w+1}^t y_{h,\tau}^w, \quad \chi_h^w \leq \sum_{t \in \mathcal{T}} y_{h,t}^w \leq 1, \quad (17)$$

$$y_{h,t}^w = 0 \quad \forall t \notin [t_h^{e,w}, t_h^{l,w} - D_h^w + 1], \quad P_{h,t}^w = \bar{P}_h^w x_{h,t}^w, \quad (18)$$

for all $w \in \mathcal{W}$, $h \in \mathcal{H}$ and $t \in \mathcal{T}$. The binary start indicator $y_{h,t}^w$ marks the step at which the cycle begins, and the convolution in (17) sets the running indicator $x_{h,t}^w$ for the following D_h^w steps. Whether a cycle is *required* on the simulated day is a per-home Bernoulli draw $\chi_h^w \in \{0, 1\}$, with $\Pr(\chi = 1) = 0.5$ for the washer/dryer and 0.6 for the dishwasher, which avoids artificial fleet-wide coincidence; where it is not required, at most one cycle may still be run if the household finds it worthwhile. The window restriction in (18) ensures completion before the deadline. The flexibility of each appliance lies in the choice of start time within its window.

2.3.8 Thermostatic loads (heat pump and refrigerator)

The dwelling, conditioned by a reversible air-source HP, and the refrigerator cabinet are both represented by a first-order (1R1C) thermal model discretised in exponential form. For $x \in \{\text{HP}, \text{fr}\}$, with temperature $T_{h,t}^x$ (indoor air or cabinet), ambient $T_{h,t}^{x,a}$ (outdoor air T_t^{amb} or kitchen air $T_{h,t}^{\text{kit}}$), and heating and cooling duty fractions $\phi_{h,t}^{x,+}$ and $\phi_{h,t}^{x,-}$,

$$T_{h,t}^x = \alpha_h^x T_{h,t-1}^x + (1 - \alpha_h^x) \left(T_{h,t}^{x,a} + R_h^x \text{COP}_h^x \bar{P}_h^x (\phi_{h,t}^{x,+} - \phi_{h,t}^{x,-}) \right), \quad \alpha_h^x = e^{-\Delta t / (R_h^x C_h^x)}, \quad (19)$$

$$\underline{T}_h^x \leq T_{h,t}^x \leq \bar{T}_h^x, \quad (20)$$

$$P_{h,t}^x = \bar{P}_h^x (\phi_{h,t}^{x,+} + \phi_{h,t}^{x,-}), \quad \phi_{h,t}^{x,+} + \phi_{h,t}^{x,-} \leq 1, \quad \phi_{h,t}^{x,\pm} \in [0, 1], \quad (21)$$

with $\phi_{h,t}^{\text{fr},+} \equiv 0$ because the refrigerator only cools. The temperature relaxes toward the steady state corresponding to the current duty at the retention factor α_h^x ; the duties enter the thermal recursion through their difference and the electrical demand through their sum, and their sum constraint forbids simultaneous heating and cooling. Representing the compressor duty as a continuous fraction rather than an on/off cycle keeps the model linear.

The HP parameters are representative of an air-source unit serving a typical dwelling: thermal resistance R_h^{HP} drawn uniformly from 2.2–3.0 °C/kW and capacitance C_h^{HP} from 8.5–11.5 kWh/°C, giving time constants of 20–33 h (mean 26 h), COP_h^{HP} drawn uniformly over 2.91–3.40 (mean 3.15), and comfort bands whose lower limit is drawn from 19.5–20.5 °C and whose width is drawn from 1.8–2.4 °C (overall 19.5–22.8 °C). On the summer study day (outdoor temperature 18–28 °C), the unit operates predominantly in cooling mode. The building's thermal inertia allows the HEMS to pre-cool toward the lower edge of the comfort band before a constrained period. Because the thermal store loses heat to a warmer ambient, pre-cooling returns less energy than it absorbs; the implications are quantified in Section 3.7. The refrigerator is confined to the food-safety band [2, 8] °C and, with an ≈ 8 h time constant, offers a small amount of short-horizon flexibility.

Layer 1 therefore returns $\mathbf{P}_h^{\text{CO}}$ and $\mathbf{P}_h^{\text{EM}}$, together with the flexibility profiles of Section 2.4, on the first pass and the re-dispatched schedule of (3) on the second. Layer 2 turns the first pair into hard operating envelopes with import and export limits, and Layer 3 sizes and verifies the network relief.

2.4 Layer 2: Aggregator envelopes, payback and settlement

The aggregator turns the Layer-1 schedules and flexibility profiles into the hard bounds that Layer 1 must respect on re-dispatch, in four steps developed below: it aggregates each home's flexibility, splits the network relief request across homes, returns the displaced energy through an energy-bounded payback, and settles the result so that no participant is worse off.

2.4.1 State-aware import and export flexibility

For each time step, the flexibility available to each home is quantified *relative to its own cost-optimal schedule* $\mathbf{P}_h^{\text{CO}}$, by reading the device powers, states of charge and temperatures along the CO trajectory. The computation uses only the home's own schedules and device states, so it can be performed locally by each HEMS immediately after the CO and EM solves, with only the three resulting connection-point power profiles uploaded to the aggregator (Table 1); in the simulations it is executed in the aggregator step from each home's own CO and EM results, which is computationally equivalent. Three quantities are produced per home, the enforceable import-reduction capacity $f_{h,t}^{\text{imp}}$, the enforceable export-reduction capacity $f_{h,t}^{\text{exp}}$ and the voluntary

recovery headroom $f_{h,t}^{\text{rec}}$:

$$f_{h,t}^{\text{imp}} = \min\left(\sum_{a \in \mathcal{D}} P_{h,t}^{a,\text{CO}} + \sum_s \Delta P_{h,t}^{\text{dis},s}, P_{h,t}^{\text{imp},\text{CO}}, (P_{h,t}^{\text{imp},\text{CO}} - P_{h,t}^{\text{imp},\text{EM}})^+\right), \quad (22a)$$

$$f_{h,t}^{\text{exp}} = P_{h,t}^{\text{exp},\text{CO}}, \quad (22b)$$

$$f_{h,t}^{\text{rec}} = \sum_s \Delta P_{h,t}^{\text{ch},s} + \Delta P_{h,t}^{\text{EWH},\uparrow} + \Delta P_{h,t}^{\text{HP},\uparrow}, \quad (22c)$$

where $(\cdot)^+ = \max(\cdot, 0)$, $\mathcal{D} = \{\text{EWH}, \text{HP}, \text{W/D}, \text{DW}, \text{EV-ch}, \text{BESS-ch}\}$ indexes the sheddable or shiftable device actions, whose CO powers $P_{h,t}^{a,\text{CO}}$ could in principle be deferred, and the storage and thermal headrooms are

$$\begin{aligned} \Delta P_{h,t}^{\text{dis},s} &= u_{h,t}^s \min\left(\bar{P}_h^{\text{dis},s} - P_{h,t}^{\text{dis},s,\text{CO}}, \frac{E_{h,t}^{s,\text{CO}} - \underline{\sigma}^s \bar{E}_h^s}{\Delta t}\right)^+, \\ \Delta P_{h,t}^{\text{ch},s} &= u_{h,t}^s \min\left(\bar{P}_h^{\text{ch},s} - P_{h,t}^{\text{ch},s,\text{CO}}, \frac{\bar{E}_h^s - E_{h,t}^{s,\text{CO}}}{\eta_{h,t}^{\text{ch},s} \Delta t}\right)^+, \\ \Delta P_{h,t}^{\text{EWH},\uparrow} &= (Q_h^{\text{EWH}} - P_{h,t}^{\text{EWH},\text{CO}}) \mathbf{1}[T_{h,t}^{\text{EWH},\text{CO}} < \bar{T}_h^{\text{EWH}} - 0.5], \\ \Delta P_{h,t}^{\text{HP},\uparrow} &= (\bar{P}_h^{\text{HP}} - P_{h,t}^{\text{HP},\text{CO}}) \mathbf{1}[T_{h,t}^{\text{HP},\text{CO}} > \underline{T}_h^{\text{HP}} + 0.25], \end{aligned} \quad (23)$$

with temperatures in $^\circ\text{C}$ and $\mathbf{1}[\cdot]$ the indicator function. Each step is evaluated independently along the CO trajectory, and the device-level sum in (22a) ignores comfort, cycle and inter-temporal energy limits, so on its own it is an optimistic screen rather than a deliverable quantity. The enforceable capacity is set by the two clips that follow it: the scheduled import and, crucially, the gap between the cost-optimal and the energy-minimal schedules. Because the EM schedule satisfies every device, comfort and cycle constraint, this clip bounds the offer by a dispatch that the home can actually execute: the resulting cap cannot fall below $P_{h,t}^{\text{imp},\text{EM}}$, and the curtailed dispatch remains feasible. The export-reduction capacity (22b) is the home's own scheduled export. It is offered to the DSO unclipped; the part of an allocated export cut that would take the export cap below the floor $P_{h,t}^{\text{exp},\text{EM}}$ of (28b) is not enforced, and any violation that results is caught by the AC-verified loop (Section 2.5.2). The recovery headroom (22c) is offered only to the payback stage and is never sold to the network as relief. The aggregator sums the uploaded per-home quantities to a node bundle,

$$F_{k,p,t}^{\text{imp}} = \sum_{h \in \mathcal{H}_{k,p}} f_{h,t}^{\text{imp}}, \quad F_{k,p,t}^{\text{exp}} = \sum_{h \in \mathcal{H}_{k,p}} f_{h,t}^{\text{exp}}, \quad (24)$$

which is the flexibility the aggregator can offer the DSO at node (k, p) . Carried alongside are the *feasibility floors* from the energy-minimal dispatch $\mathbf{P}_h^{\text{EM}}$: its import profile $P_{h,t}^{\text{imp},\text{EM}}$ and the PV surplus $P_{h,t}^{\text{exp},\text{EM}}$ it cannot absorb locally. These floors anchor the hard caps derived below and guarantee their feasibility.

2.4.2 Node request split and energy-bounded payback

Given the network relief request from Layer 3 (Section 2.5), the aggregator assembles the per-home hard bounds. The node-level import-cut and export-cut requests returned by Layer 3 are split among the homes of a bus-phase node in proportion to their enforceable capacities $f_{h,t}^{\text{imp}}$ and $f_{h,t}^{\text{exp}}$, yielding the per-home *allocated cuts* $\delta_{h,t}^{\text{imp}} \leq f_{h,t}^{\text{imp}}$ and $\delta_{h,t}^{\text{exp}} \leq f_{h,t}^{\text{exp}}$; the import cut is additionally re-clipped by (22a). This proportional split allocates the burden to the homes best able to bear it. Because curtailment shifts *energy* rather than merely clipping power, the deferred energy must be assigned a period in which it can be recovered; otherwise, it rebounds and may create a new peak. A single linear program (LP) therefore grants recovery headroom

$r_{h,t} \geq 0$ at un-curtailed steps:

$$\mathbf{r}^{(1)}, \mathbf{r}^{(2)} \geq 0, \pi \geq 0 \quad \sum_{h \in \mathcal{H}} \sum_{t \in \mathcal{T}} \left(w^{\text{need}} \omega_h r_{h,t}^{(1)} + w^{\text{extra}} r_{h,t}^{(2)} \right) \Delta t - w^{\text{pk}} \sum_{n \in \mathcal{F}} \pi_n \quad (25a)$$

$$\text{s.t.} \quad r_{h,t} = r_{h,t}^{(1)} + r_{h,t}^{(2)} \leq f_{h,t}^{\text{rec}}, \quad \forall h, t, \quad (25b)$$

$$\sum_{t \in \mathcal{T}} r_{h,t}^{(1)} \Delta t \leq \eta^{\text{rec}} E_h^{\text{cut}}, \quad \forall h, \quad (25c)$$

$$\sum_{h \in \mathcal{H}} (P_{h,t}^{\text{cut}} + r_{h,t}) \leq \text{PF} \bar{S}_{\text{tx}}^{\text{loss}}, \quad \forall t, \quad (25d)$$

$$\sum_{h \in n} (P_{h,t}^{\text{cut}} + r_{h,t}) \leq \pi_n, \quad \forall n, t, \quad (25e)$$

$$\sum_{h \in n} (P_{h,t}^{\text{cut}} + r_{h,t}) \leq \bar{P}_n^{\text{base}}, \quad \forall n, t, \quad (25f)$$

in which the post-curtailement net power of each home is

$$P_{h,t}^{\text{cut}} = P_{h,t}^{\text{imp,CO}} - P_{h,t}^{\text{exp,CO}} - \delta_{h,t}^{\text{imp}} + \delta_{h,t}^{\text{exp}}. \quad (26)$$

Feeders are indexed by $n \in \mathcal{F}$ and $\pi_n \geq 0$ is the peak variable of feeder n . Recovery is granted in two tiers: the tier-one variable $r_{h,t}^{(1)}$ returns a home its *own* curtailed energy E_h^{cut} up to the recoverable fraction η^{rec} imposed by (25c), and the tier-two variable $r_{h,t}^{(2)}$ may then distribute any remaining headroom. The study uses $w^{\text{need}}=10$ and $w^{\text{extra}}=0$, so only tier-one recovery is rewarded; this prevents the program from granting import headroom that is not justified by deferred energy and would otherwise be absorbed as additional consumption. The per-home coefficient

$$\omega_h = 1 + \gamma \frac{\vartheta_h}{\bar{\vartheta}}, \quad \vartheta_h = \frac{\sum_{t \in \mathcal{T}} (\delta_{h,t}^{\text{imp}} + \delta_{h,t}^{\text{exp}}) \Delta t}{E_h^{\text{shift}}}, \quad \bar{\vartheta} = \frac{1}{|\mathcal{H}|} \sum_{j \in \mathcal{H}} \vartheta_j, \quad (27)$$

grades tier-one recovery by each household's *curtailment burden ratio* ϑ_h , normalised by the fleet mean $\bar{\vartheta}$, at fairness strength γ . The burden ratio relates the energy a home is asked to give up to its total shiftable energy E_h^{shift} , so that large homes are not scored as heavily burdened merely because of their size. The study uses $\gamma = 0$ ($\omega_h \equiv 1$), so (27) is an available option rather than a mechanism on which the results depend. Constraint (25b) keeps the grant within the voluntary headroom, and (25d) limits the fleet's post-curtailement power plus recovery to the loss-adjusted substation rating $\bar{S}_{\text{tx}}^{\text{loss}}$ (the nameplate rating net of a loss allowance) at power factor PF. Constraint (25e) penalises the feeder peak so that payback fills valleys, and (25f) ensures that recovery never raises a feeder above its uncoordinated maximum $\bar{P}_n^{\text{base}}$. Since (25c) is an inequality and $r_{h,t}$ enters (28a) as an addition to the import cap, payback is a bounded recovery *opportunity* rather than an obligation: a household uses the headroom only to the extent that its devices have deferred energy to recover. Compelling a home to import its full grant would instead force energy into the HP load, where it would dissipate within the horizon. The LP is always feasible, since $\mathbf{r} = 0$ is admissible.

2.4.3 Hard import and export limits with feasibility floors

The envelope broadcast to each home is a pair of hard bounds on net power $P_{h,t}^{\text{net}} = P_{h,t}^{\text{imp}} - P_{h,t}^{\text{exp}}$,

$$\bar{P}_{h,t}^{\text{DOE}} = \max \left(P_{h,t}^{\text{imp,CO}} - \delta_{h,t}^{\text{imp}} + r_{h,t}, P_{h,t}^{\text{imp,EM}}, 0 \right) + \varepsilon, \quad (28a)$$

$$-P_{h,t}^{\text{DOE}} = \max \left(P_{h,t}^{\text{exp,CO}} - \delta_{h,t}^{\text{exp}}, P_{h,t}^{\text{exp,EM}}, 0 \right). \quad (28b)$$

Equation (28a) defines the import cap $\bar{P}_{h,t}^{\text{DOE}}$ and (28b) the export cap $-P_{h,t}^{\text{DOE}}$; $\varepsilon \geq 0$ is an optional numerical margin on the import cap, set to zero in this study, in which no re-dispatch was made infeasible by solver tolerances. The floors constitute the essential design element: $\bar{P}_{h,t}^{\text{DOE}} \geq P_{h,t}^{\text{imp,EM}}$ and the export cap $\geq P_{h,t}^{\text{exp,EM}}$ guarantee that the energy-minimal dispatch remains admissible, so the Layer-1 re-dispatch MILP under these *hard* caps is feasible without any envelope slack: compliance is a guarantee, not an incentive. This property is stated formally below.

Proposition 1 (Feasibility of the re-dispatch). For every home h , the energy-minimal schedule $\mathbf{P}_h^{\text{EM}}$ satisfies all constraints of the re-dispatch problem (3) with $s_{h,t}^{\text{DOE}} = 0$ for all $t \in \mathcal{T}$, for any allocated cuts $\delta_{h,t}^{\text{imp}}$, $\delta_{h,t}^{\text{exp}}$ and payback grants $r_{h,t} \geq 0$.

Proof. The re-dispatch shares the device and grid constraints (5)–(21) with the EM solve, so $\mathbf{P}_h^{\text{EM}}$ satisfies them. By the complementarity constraint, at each step $\mathbf{P}_h^{\text{EM}}$ either imports $P_{h,t}^{\text{imp,EM}}$ or exports $P_{h,t}^{\text{exp,EM}}$, so its net power lies in $[-P_{h,t}^{\text{exp,EM}}, P_{h,t}^{\text{imp,EM}}]$. By (28), $\bar{P}_{h,t}^{\text{DOE}} \geq P_{h,t}^{\text{imp,EM}}$ and $-\bar{P}_{h,t}^{\text{DOE}} \geq P_{h,t}^{\text{exp,EM}}$, so (8) holds with zero slack. Because the tightening rule (38) never moves a cap below its floor, the property holds at every iteration of the outer loop. $\square$

Because the proposed scheme fixes the slack at zero, Proposition 1 is what ensures that every re-dispatch problem is feasible; none was infeasible in the study.

2.4.4 Settlement and bill protection

Once the dispatch is known, an ex-post settlement pays each household a per-kWh service fee λ^{fee} on its import reduction relative to its CO schedule, capped at its allocated cut $\delta_{h,t}^{\text{imp}}$, and, where the fee is insufficient, adds a bill-protection credit that closes the remaining gap, enforcing

$$B_h \leq B_h^{\text{CO}} \quad \forall h, \quad (29)$$

so that no household is worse off than under its uncoordinated schedule. The fee is paid on the import side only; export reduction, which constitutes the larger share of the relief on this feeder, is compensated implicitly, since the forgone revenue enters B_h and is captured by the bill-protection term. The fee should therefore not be interpreted as the price of the network service; the guarantee rests on each household's own counterfactual bill rather than on a posted price. The settlement is a transfer funded by the network service and chargeable against the reinforcement the DSO defers; the regulated mechanism for this cost recovery is not modelled. Because households that gain are not charged, the total transfer (service fees plus bill protection) is at least the sum over homes of the pre-settlement bill increase relative to B_h^{CO} , i.e. $|\mathcal{H}|$ times the fleet-mean gross premium of Section 3.9.

2.5 Layer 3: DSO AC-OPF request and AC-verified tightening

2.5.1 Minimum-deployment flexibility request

The DSO sizes the network relief by solving an unbalanced three-phase AC optimal power flow over the full horizon. Distribution feeders are inherently unbalanced, so a single-phase equivalent is inadequate; the rectangular current-mismatch formulation of [63] is therefore adopted, which expresses each nodal voltage by its real and imaginary parts and enforces power balance through the injected current. Buses are indexed by $k \in \mathcal{N}$, phases by $p \in \mathcal{P} = \{a, b, c\}$ and the substation by $k = s$. The objective deploys the *least* flexibility that clears every violation:

$$\min_{\mathbf{P}^{\text{flex}}} w_F \sum_{k,p,t} |P_{k,p,t}^{\text{flex}}| + w_L \sum_t P_{\text{loss},t} + w_S \sum_{c \in \mathcal{C}} s_c, \quad (30)$$

where w_F , w_L and w_S weight the total flexibility deployed, the network losses, and a set of non-negative violation slacks s_c indexed by $c \in \mathcal{C}$. The slack set spans every softened network limit, namely the upper and lower voltage bounds, the line-current limit and the transformer apparent-power limit, all penalised at the same weight. The first term makes the request *minimum-deployment*; the slack term is a soft penalty that lets the solver expose an infeasible penetration level, on any of the three limit families, rather than fail outright. The optimisation is subject to the unbalanced three-phase network equations, written in per unit. Let $V_{k,p,t} = V_{k,p,t}^{\text{re}} + jV_{k,p,t}^{\text{im}}$ be the phase- p voltage at bus k , $\mathbf{V}_t$ and $\mathbf{I}_t$ the vectors of all nodal voltages and current injections at step t , and $\mathbf{Y}$ the three-phase bus admittance matrix. For every bus, phase and time step,

$$\mathbf{I}_t = \mathbf{Y}\mathbf{V}_t, \quad (31)$$

$$(P_{k,p,t}^g + P_{k,p,t}^{\text{flex}} + P_{k,p,t}^{\text{exp}} - P_{k,p,t}^d) + j(Q_{k,p,t}^g + Q_{k,p,t}^{\text{flex}} - Q_{k,p,t}^d) = V_{k,p,t} I_{k,p,t}^*, \quad (32)$$

in which $P_{k,p,t}^g, Q_{k,p,t}^g$ are the substation injections, $P_{k,p,t}^d, Q_{k,p,t}^d$ the aggregated household demand at the node, $P_{k,p,t}^{\text{exp}}$ the aggregated export, and $P_{k,p,t}^{\text{flex}}, Q_{k,p,t}^{\text{flex}}$ the flexibility deployed. Operating limits are imposed on voltage magnitude, line current and transformer apparent power, and the substation voltage is fixed at its reference:

$$\underline{V}^2 \leq |V_{k,p,t}|^2 \leq \bar{V}^2, \quad (33)$$

$$|I_{\ell,p,t}|^2 \leq (\kappa \bar{I}_{\ell})^2, \quad \mathbf{I}_{\ell,t} = \mathbf{Y}_{k(\ell)k'(\ell)} (\mathbf{V}_{k'(\ell),t} - \mathbf{V}_{k(\ell),t}), \quad (34)$$

$$\sum_{p \in \mathcal{P}} |S_{s,p,t}^g|^2 \leq \bar{S}_{\text{tx}}^2, \quad S_{s,p,t}^g = P_{s,p,t}^g + jQ_{s,p,t}^g, \quad (35)$$

$$V_{s,p,t} = |V^{\text{ref}}| e^{j\theta_p}. \quad (36)$$

In (34), $\mathbf{I}_{\ell,t}$ collects the three phase *branch* currents of line ℓ , distinct from the nodal injections of (31), $k(\ell)$ and $k'(\ell)$ are its sending and receiving buses, $\mathbf{Y}_{kk'}$ is the corresponding 3×3 off-diagonal block of $\mathbf{Y}$, and the ampacity $\bar{I}_\ell$ is expressed in per unit. The model is implemented in rectangular coordinates, so that (31) is linear and (32)–(35) are quadratic in $(V^{\text{re}}, V^{\text{im}})$. In (36) the substation is held at its nominal magnitude with the balanced three-phase reference angles $\theta_a = 0^\circ$, $\theta_b = -120^\circ$, $\theta_c = +120^\circ$. The substation injections exist only at the slack bus and are identically zero for every $k \neq s$; at $k = s$ the balance equations (32) therefore serve only to determine the substation injection, whose voltage is fixed by (36), while (31) holds at every bus.

The transformer constraint (35) is imposed on the per-phase apparent powers of the branches leaving the substation secondary. Summing the squared per-phase apparent powers, rather than squaring the three-phase totals, makes (35) a looser bound than a strict three-phase rating by a factor of up to $\sqrt{3}$; the substation peaks at 0.30 p.u. of rating in the worst case tested (Section 3.5), so the constraint is never active and the distinction does not affect any reported result. Each of (33)–(35) carries its own slack from $\mathcal{C}$, omitted above for readability. The factor κ in (34) sets the fraction of ampacity against which the request is sized. The proposed framework uses $\kappa = 1.0$, so the request targets the same limit against which the OpenDSS replay is judged (Section 2.6). With $\kappa < 1$, the OPF would size relief for loading that the acceptance test regards as compliant (for example, the proposed case's 1.02–1.03 p.u. overnight charging peaks at 75% and 100% DER (Section 3.5), which exceed the 1.00 p.u. planning limit but remain below nameplate ampacity), and the curtailed energy would rebound into previously violation-free slots. The security margin is instead provided by the AC-verified replay, which is authoritative, whereas the request is a model prediction. The conventional benchmark retains the usual $\kappa = 0.9$ planning margin.

Independently of κ , all line loadings are reported on a common per-unit base of 90% of conductor ampacity, so that 1.00 p.u. is the planning limit and 1.11 p.u. is nameplate ampacity, the violation threshold. Finally, the flexibility deployed at each node is bounded by the aggregator's enforceable envelope from (24), mapping the per-home limits of (22) to the node,

$$-F_{k,p,t}^{\text{exp}} \leq P_{k,p,t}^{\text{flex}} \leq F_{k,p,t}^{\text{imp}}, \quad \forall k \in \mathcal{N}, p \in \mathcal{P}, t \in \mathcal{T}, \quad (37)$$

and the aggregated reactive demand and reactive flexibility follow the load at a constant power factor (PF), $Q_{k,p,t}^\bullet = P_{k,p,t}^\bullet \tan(\arccos \text{PF})$. The study uses $\text{PF} = 1.0$, so $Q^d = Q^{\text{flex}} = 0$. The absolute value in (30) is not evaluated directly: each $|P_{k,p,t}^{\text{flex}}|$ is replaced by an auxiliary variable $\zeta_{k,p,t} \geq 0$ constrained by $P_{k,p,t}^{\text{flex}} \leq \zeta_{k,p,t}$ and $-P_{k,p,t}^{\text{flex}} \leq \zeta_{k,p,t}$, which is exact at the optimum because ζ is minimised. The program (30)–(37) is a non-convex nonlinear program, and is solved with the interior-point method IPOPT [64], which certifies local optimality only. Each solve is warm-started from the OpenDSS operating point, and the locally optimal request is independently verified by the replay of Section 2.5.2. The per-node solution $\{P_{k,p,t}^{\text{flex}}\}$ is the relief request that Layer 2 converts to hard caps. The AC-OPF is solved once per day-ahead evaluation, and the tightening loop iterates on the OpenDSS replay alone.

2.5.2 AC-verified tightening loop

Because the AC-OPF is an approximation of the operating network, the realised HEMS profiles are replayed through a full OpenDSS power flow, from which bus voltages and line currents are read directly; transformer loading is evaluated from the loss-adjusted net power at the feeder head, which is adequate because the substation is never close to binding. Any residual violation tightens the envelopes *only in the time steps at which it occurs*, geometrically toward the feasibility floors,

$$\bar{P}_{h,t}^{(i+1)} = \underline{P}_{h,t}^{\text{fl}} + \beta (\bar{P}_{h,t}^{(i)} - \underline{P}_{h,t}^{\text{fl}}), \quad \beta = 0.6, \quad i = 0, \dots, K_{\max} - 2, \quad (38)$$

where $\bar{P}_{h,t}^{(i)}$ is the magnitude of the cap being tightened (the import cap $\bar{P}_{h,t}^{\text{DOE}}$ or the export cap $-\underline{P}_{h,t}^{\text{DOE}}$) and $\underline{P}_{h,t}^{\text{fl}}$ its feasibility floor ($P_{h,t}^{\text{imp,EM}}$ or $P_{h,t}^{\text{exp,EM}}$). The rule shrinks export caps at over-voltage/reverse-overload slots and import caps at under-voltage/forward-overload slots (transformer overload tightens import caps network-wide). The tightening is *temporal rather than spatial*: a violation anywhere on a feeder tightens the corresponding cap of every home on that feeder, but only in the half-hour in which it occurs. This is justified because an LV feeder is electrically short, so nearly all homes contribute to a given conductor current, and because per-node attribution would require the aggregator to hold a network sensitivity model, which principle (P1) assigns to the DSO. The drawback is that non-contributing homes are also tightened; the benefit is that temporal selectivity, where the binding events are concentrated, requires no additional information. The loop halts on the first violation-free replay or when the iteration budget $K_{\max}$ is exhausted, in which case the residual violation is reported as the feeder's hosting-capacity boundary; convergence is discussed in Section 2.12.

2.6 Hosting-capacity metric

For the penetration sweep, the feeder hosting capacity is the highest penetration level (in per cent of the 55 customers) at which the full 24 h simulation produces no network violation. A violation is declared when bus voltage leaves the 0.90–1.10 p.u. band of EN 50160, when line loading exceeds nameplate ampacity (1.11 p.u. on the reporting base defined above), or when the

substation transformer exceeds its short-term limit of 1.2 p.u. of rating. The voltage test is applied *strictly*, to the unrounded per-unit magnitude, so that a bus at 1.1009 p.u. counts as a breach. Because EN 50160 is stated for 95% of 10-min mean values over a week, applying it as a hard limit at every half-hourly step is conservative. Although the UK statutory limits of the Electricity Safety, Quality and Continuity Regulations (ESQCR) 2002 are 0.94–1.10 p.u., the minimum bus voltage never falls below 0.98 p.u., so all results also satisfy the ESQCR floor. The short-term transformer threshold is used because distribution transformers tolerate brief overloads without significant loss of life. The resulting metric is an *operational* hosting capacity determined by the tightening loop rather than pre-computed.

2.7 Framework architecture and design principles

The models above are assembled into the coordination architecture of Fig. 1. The three layers are organised by *actor* rather than by time scale: the home (Layer 1), the aggregator (Layer 2) and the DSO (Layer 3) each solve their own optimisation over the same day-ahead horizon and exchange signals between solves. This decomposition of authority places the quantities that must remain private (device states, occupancy and comfort preferences) and those that must be shared (net power at the connection point) on opposite sides of the same organisational boundary. Four principles govern the design.

(P1) *Subsidiarity*. Each decision is taken at the lowest layer that holds the required information. Device scheduling requires device states and occupant preferences and therefore remains in the home; sizing the network relief requires the admittance matrix and nodal injections and therefore remains with the DSO. The aggregator holds neither and is confined to translating a node-level relief request into per-home bounds.

(P2) *Minimal and interpretable interface*. Only a pair of power bounds per home and time step crosses a layer boundary downward, and only connection-point power profiles (schedules and flexibility) cross upward. No device telemetry leaves the dwelling, and no network model leaves the DSO. Unlike a price signal or a direct load-control command, the envelope states what the connection point may draw and inject and leaves the means to the household.

(P3) *Feasibility by construction*. A bound that a household cannot meet results in either a violation or an infeasible re-dispatch. The feasibility floors of (28) therefore anchor every cap on the home’s maximum-self-consumption dispatch, so that a compliant re-dispatch exists before the envelope is broadcast. The tracking term of (3) selects among compliant schedules but does not determine whether one exists; the proposed scheme sets the compliance rate to zero and compensates households ex post, so compliance rests on the caps alone.

(P4) *Independent verification*. Because the relief request is sized on a model, the realised profiles are replayed through an independent unbalanced power-flow solution, and the envelopes are tightened until the replay is violation-free. Compliance is never certified by the model that sized the request.

Information flows upward as HEMS profiles and residual violations, and control signals flow downward as envelope bounds and, when the outer loop is active, as tightened bounds at the violated time slots (Fig. 1). The lower panel of the figure places the proposed scheme beside the two benchmark controllers of Section 2.13.

2.8 Information exchange and privacy

Table 1 details the interface of principle (P2). Over one day-ahead cycle, each household uploads two net power profiles and three connection-point flexibility profiles, receives a pair of bound profiles and returns its re-dispatched net profile. Per-device variables, including the state trajectories from which the flexibility profiles are computed, remain within the HEMS and are never transmitted, whereas a central scheduler would require the state of every appliance in every dwelling.

Table 1. Interface between layers over one day-ahead coordination cycle.

Interface	Direction	Payload
Home → Aggregator	up	$\mathbf{P}_h^{\text{CO}}, \mathbf{P}_h^{\text{EM}}, f_{h,t}^{\text{imp}}, f_{h,t}^{\text{exp}}, f_{h,t}^{\text{rec}}$
Aggregator → DSO	up	$F_{k,p,t}^{\text{imp}}, F_{k,p,t}^{\text{exp}}$
DSO → Aggregator	down	$P_{k,p,t}^{\text{flex}}$
Aggregator → Home	down	$\underline{P}_{h,t}^{\text{DOE}}, \bar{P}_{h,t}^{\text{DOE}}$
Home → Aggregator	up	re-dispatched $P_{h,t}^{\text{net}}$
DSO (internal)	—	OpenDSS replay, tightening

2.9 Operation of Layer 1: day-ahead scheduling

In the first pass, each prosumer independently solves the MILP of (5)–(21) under objectives (1) and (2), returning the cost-optimal schedule $\mathbf{P}_h^{\text{CO}}$ and the energy-minimal schedule $\mathbf{P}_h^{\text{EM}}$, and evaluates its flexibility profiles (22) locally along the CO state trajectory. The pair brackets the household’s admissible operating range: CO represents its preferred schedule and EM the




Figure 1. Three-layer hierarchical DER coordination framework. Layer 1 (home) computes the cost-optimal (CO) and energy-minimal (EM) schedules and re-dispatches (RD) under the received envelope. Layer 2 (aggregator) separates enforceable import- and export-reduction capacity from voluntary recovery headroom and issues hard import and export limits anchored on the EM floors, with energy-bounded payback and ex-post settlement. Layer 3 (DSO) sizes the request with an unbalanced AC-OPF and verifies it in OpenDSS, tightening caps only in violated time slots (dashed loop). The lower panel lists the three controllers compared in Section 2.13.

maximum concession it can physically provide, and every envelope issued by the aggregator is bounded below by the EM floors. In the re-dispatch pass, the household solves (3) under the hard caps (28). Since each cap dominates the corresponding EM value, the EM schedule is always feasible without envelope slack; the tracking term only selects among compliant schedules. The $|\mathcal{H}|$ problems are independent and are solved concurrently with CPLEX.

2.10 Operation of Layer 2: envelope construction

The aggregator performs four functions in sequence: (a) it aggregates the uploaded state-aware import and export flexibility profiles (22) into the node bundle (24); (b) it splits the Layer-3 relief request among the homes of each node in proportion to their enforceable capacities; (c) it solves the payback LP (25), returning the displaced energy in network-safe periods; and (d) it assembles the hard caps (28) and broadcasts them to Layer 1. The ex-post settlement (29) then keeps every household at or below its uncoordinated bill. These functions are assigned to the aggregator because it is the only actor that observes a population of homes without observing the interior of any of them.

2.11 Operation of Layer 3: relief sizing and validation

The DSO layer performs two distinct tasks. It first *sizes* the relief by solving the minimum-deployment AC-OPF (30), and then *verifies* the outcome by replaying the re-dispatched profiles through an unbalanced OpenDSS power flow, applying the tightening rule (38) at any violated time step. The two tasks are kept separate because they fail in different ways: sizing is an optimisation over an approximate model, whereas verification simulates the realised schedule and is authoritative on feeder security. Both are posed against the same current limit ($\kappa = 1.0$), so the margin against model error is supplied by the outer loop, and only where a violation is observed, rather than by a static de-rating. Unlike a last-resort escalation, the AC-OPF is part of every coordination cycle, so the request is always minimum-deployment, and the device-level dispatch that meets it remains in the home.

2.12 Operational sequence and convergence

One coordination cycle proceeds as set out in Algorithm 1. Only the re-dispatch and verification steps lie inside the outer loop. Settlement is applied to the final dispatched schedule, since only that schedule determines the curtailment for which each household is compensated. Because the AC-OPF is solved once per cycle, each additional iteration requires only a set of parallel MILPs and a single power flow, which keeps the cost of the loop low.

Termination. Applied in a violated time slot, the rule (38) gives $\bar{P}_{h,t}^{(i)} - \underline{P}_{h,t}^{\text{fl}} = \beta^i (\bar{P}_{h,t}^{(0)} - \underline{P}_{h,t}^{\text{fl}})$ with $\beta \in (0, 1)$, so every tightened cap decreases monotonically toward a feasible floor that it never crosses, while caps at non-violated slots are unchanged. The caps therefore can neither cycle nor render the re-dispatch infeasible. Since the contraction is geometric, termination is enforced by the budget $K_{\max}$, the maximum number of dispatch–replay passes (at most $K_{\max} - 1$ tightenings): the loop halts either on a violation-free replay, in which case the schedule is dispatched, or on exhaustion of the budget, in which case the residual violation is reported as the feeder’s hosting-capacity boundary under comfort-preserving flexibility. In the study of Section 3, at most one tightening pass is required.

2.13 Controller scenarios

Three controllers are compared. The comparison of Case 1 with Case 2 shows the value of network-aware coordination, and that of Case 2 with Case 3 shows the combined value of the proposed elements, whose individual contributions are isolated by the ablation of Section 3.11.

- **Decentralised baseline** (Case 1). Layer-1 HEMS only, with no aggregator, no envelopes and no network optimisation, each household minimising its own bill under the time-of-use tariff. This is the network-unaware reference of Section 1.1; because it is exactly the CO schedule of (1), it is referred to interchangeably below as the *decentralised* case and the *uncoordinated* baseline.
- **Conventional hierarchical** (Case 2). A Layer-2 aggregator issues flexibility limits on the import side only, with flexibility payments, and leaves export unconstrained, so reverse power flow is not a controlled quantity in this scheme. It is stated formally in Section 2.13.1 below. This case represents the import-reduction schemes that dominate the hierarchical HEMS-coordination literature reviewed in Section 1.2.
- **Proposed** (Case 3). The full three-layer framework: HEMS scheduling, operating envelopes with enforced import and export limits, energy-bounded payback and hard caps anchored on the EM feasibility floors, and a DSO AC-OPF request closed by the AC-verified tightening loop.

Algorithm 1 Hierarchical day-ahead coordination cycle with envelope tightening.

```
1: for each day-ahead coordination cycle do
2:   LAYER 1 (LOCAL SCHEDULING):
3:   for each home  $h \in \mathcal{H}$  do
4:     Solve CO and EM.
5:     Form  $f_{h,t}^{\text{imp}}$ ,  $f_{h,t}^{\text{exp}}$  and  $f_{h,t}^{\text{rec}}$  by (22).
6:     Upload the schedules and flexibility profiles to the aggregator.
7:   end for
8:   LAYER 2 (FLEXIBILITY BUNDLING):
9:   Aggregate the uploaded profiles to the node bundle (24).
10:  LAYER 3 (RELIEF SIZING):
11:  The DSO solves the AC-OPF (30) within the node bundle.
12:  return per-node request  $P_{k,p,t}^{\text{flex}}$ .
13:  ALLOCATION AND PAYBACK:
14:  Split  $P_{k,p,t}^{\text{flex}}$  across the homes of node  $(k, p)$ .
15:  Solve the payback LP (25).
16:  Assemble the hard caps  $\bar{P}_{h,t}^{(0)}$  by (28).
17:  OUTER LOOP (RE-DISPATCH, VERIFICATION, TIGHTENING):
18:   $i \leftarrow 0$ .
19:  RE-DISPATCH: each home re-solves (3) under  $\bar{P}_{h,t}^{(0)}$ , in parallel.
20:  VERIFICATION: replay the realised profiles in OpenDSS.
21:  while the replay is not violation-free and  $i < K_{\max} - 1$  do
22:    TIGHTENING: at each violated slot  $t$ , for every home on the affected feeder,
23:     $\bar{P}_{h,t}^{(i+1)} - \underline{P}_{h,t}^{\text{fl}} = \beta (\bar{P}_{h,t}^{(i)} - \underline{P}_{h,t}^{\text{fl}})$ ,  $\beta \in (0, 1)$ , by (38);
24:    caps at non-violated slots are left unchanged.
25:     $i \leftarrow i + 1$ .
26:    Re-dispatch under  $\bar{P}_{h,t}^{(i)}$  and replay in OpenDSS.
27:  end while
28:  if the replay is violation-free then
29:    Dispatch the schedule.
30:    SETTLEMENT: apply the ex-post settlement (29)
31:    to the dispatched schedule.
32:  else
33:    Report the residual violation as the feeder's hosting-capacity
34:    boundary under comfort-preserving flexibility.
35:  end if
36: end for
```

2.13.1 Formal statement of the conventional benchmark

Because the central comparison of this paper is against Case 2, its formulation is stated explicitly. It differs from the proposed scheme in four places, which are chosen to reproduce the import-reduction coordination designs it represents [42, 43] rather than to handicap it. In the figures and the summary table this case is labelled *Hierarchical*.

(i) *Request*. Case 2's DSO layer minimises the *signed* deployment rather than its magnitude, and couples a fixed reactive component to it,

$$\min w_F^{(2)} \sum_{k,p,t} P_{k,p,t}^{\text{flex}} + w_L \sum_t P_{\text{loss},t} + w_S \sum_{c \in \mathcal{C}} s_c, \quad Q_{k,p,t}^{\text{flex}} = \psi^{\text{flex}} P_{k,p,t}^{\text{flex}}, \quad (39)$$

with $w_F^{(2)} = 1$ and $\psi^{\text{flex}} = 0.60$, against the proposed scheme's $w_F = 5$ on $|P^{\text{flex}}|$ and $\psi^{\text{flex}} = 0$. Minimising the signed sum rather than the absolute value means that a *negative* deployment, that is an instruction to increase demand, reduces Case 2's objective instead of costing it, so the benchmark is structurally inclined to request demand increases wherever its bundle admits them.

(ii) *Bundle, floor and trigger*. Case 2's node bundle is the gap between the cost-optimal import and the gross controllable load of the energy-minimal schedule, $F_{k,p,t}^{(2)} = \sum_{h \in \mathcal{H}_{k,p}} (P_{h,t}^{\text{imp,CO}} - P_{h,t}^{\text{gross,EM}})$, which is negative wherever the cost-optimal

schedule exports. The request is bounded by this bundle and, where the bundle is large enough, by the published method's minimum-deployment floor, a fixed fraction of the node's CO import $P_{k,p,t}^{\text{imp,CO}} = \sum_{h \in \mathcal{H}_{k,p}} P_{h,t}^{\text{imp,CO}}$:

$$P_{k,p,t}^{\text{flex}} \in \begin{cases} [\mu^{\text{fl}} P_{k,p,t}^{\text{imp,CO}}, F_{k,p,t}^{(2)}], & F_{k,p,t}^{(2)} > \mu^{\text{fl}} P_{k,p,t}^{\text{imp,CO}}, \\ [0, F_{k,p,t}^{(2)}], & 0 \leq F_{k,p,t}^{(2)} \leq \mu^{\text{fl}} P_{k,p,t}^{\text{imp,CO}}, \\ [\max(F_{k,p,t}^{(2)}, -P_{k,p,t}^{\text{exp,CO}}), 0], & F_{k,p,t}^{(2)} < 0, \end{cases} \quad \mu^{\text{fl}} = 0.05, \quad (40)$$

with $P_{k,p,t}^{\text{exp,CO}}$ the node's CO export. Case 2 is computed at every penetration level, but its re-dispatch is adopted only if its replay strictly reduces the total violation severity of the decentralised replay; otherwise the baseline is reported. It therefore reproduces the decentralised baseline exactly below the first violation.

(iii) *Per-home limit.* The node request is mapped to a home by a share ξ_h of the node total and becomes an import limit with no corresponding export limit,

$$\bar{P}_{h,t}^{\text{DOE,(2)}} = \begin{cases} P_{h,t}^{\text{gross,EM}} - \xi_h P_{k_h,p_h,t}^{\text{flex}}, & P_{h,t}^{\text{gross,EM}} - \xi_h P_{k_h,p_h,t}^{\text{flex}} \geq 0, \\ P_{h,t}^{\text{gross,EM}}, & \text{otherwise,} \end{cases} \quad \bar{P}_{h,t}^{\text{DOE,(2)}} = -\infty. \quad (41)$$

For reporting, the aggregator's per-home import target is $P_{h,t}^{\text{imp,CO}} - \xi_h P_{k_h,p_h,t}^{\text{flex}}$, where (k_h, p_h) is the node of home h ; it enters no objective.

(iv) *Payment.* Compliance is procured by price rather than guaranteed by construction: the home receives the rate $\lambda_{h,t}^{\text{inc}}$ of (3) on imported energy, a per-home, per-step logistic rate $\lambda_{h,t}^{\text{inc}} = -0.06 + 0.08 \sigma(P_{h,t}^{\text{gross,EM}} - P_{h,t}^{\text{imp,CO}})$ £/kWh, with $\sigma(x) = 1/(1 + e^{-x})$ and x in kWh, which approaches +£0.02/kWh where the energy-minimal gross load exceeds the cost-optimal import and -£0.06/kWh (a charge) where the cost-optimal schedule imports more, and the import limit (41) is enforced through the slack $s_{h,t}^{\text{DOE}}$ penalised at $w^{\text{env}} = 10^3$ £/kWh (the realised slack is zero). Case 2's re-dispatch has no tracking term. There is no ex-post bill-protection stage, so a household's realised bill may exceed its uncoordinated bill, and the sigmoid payment may be negative, i.e. a net charge. Both consequences are visible in the results (Section 3.9).

One caveat affects the interpretation of the comparison. The mapping (41) is defined for an import-reduction request ($P^{\text{flex}} > 0$), whereas on this feeder the AC-OPF frequently returns $P^{\text{flex}} < 0$ in the export window: the signed objective drives every node with a negative bundle to its lower bound in (40), so that at 100% DER 924 of the 2,640 node-step requests are negative (-1,052 kWh/day in total), concentrated between 08:00 and 18:30. Lacking an export cap, (41) converts such a request into a raised import ceiling. This is a faithful consequence of the published design, since an import-only aggregator can respond only by adjusting import, and it implies that Case 2's shortfall should be attributed to the absence of an enforced export limit rather than to its tuning. Section 3.11 tests this attribution directly by releasing the export bound within the proposed framework. That ablation (A1) is also the like-for-like economic reference: with every other setting of the proposed scheme retained, its gross cost is £1.07 and £0.40/day at 75% and 100% (Table 6), against £1.40 and £0.69/day for Case 2, so part of Case 2's bill premium reflects its sigmoid payment, its request design and the absence of bill protection rather than its import-only structure alone.

Because Case 3 solves (30) in every cycle, and its objective includes a loss term while every dwelling owns thermal and wet-appliance flexibility, it returns a small non-zero request even when no limit is threatened. It therefore departs slightly from the baseline at every penetration level, including 0%. In return, when substantial relief is required, it is sized against the entire horizon and feeder simultaneously rather than assembled from successive post-violation escalations. All three cases use the same network model of Section 2.2 and the identical device population of Section 3.2.

3 Results and discussion

3.1 Test network and setup

The study uses the IEEE European Low-Voltage Test Feeder [65], an unbalanced three-phase LV network (11/0.416 kV, 906 buses, 55 single-phase residential connections) solved in OpenDSS [66]; its parameters are listed in Table 2. A 24 h day-ahead simulation (48 half-hourly steps) is run for each DER penetration level $\Pi = \{0, 25, 50, 75, 100\}\%$, where $v\%$ denotes that round(55v/100) households (0, 14, 28, 41 and 55) own the full DER bundle of PV, EV and battery. All randomised household parameters use a fixed seed.

The feeder consists of a long, resistive trunk ($R/X \approx 6$) supplying short single-phase laterals, so active power is the dominant control variable for both voltage and loading. Under the tariff described below, least-cost scheduling charges vehicles and batteries in the 02:00–05:00 low-price window, tops the batteries up from midday PV, and discharges the batteries simultaneously when the export price rises at 16:00. The uncoordinated feeder is therefore subject to two stress events, of

Table 2. Test feeder and per-home device parameters. Line loading is reported per unit of the 90%-of-ampacity planning limit against which the DSO sizes its request, so that 1.00 p.u. is the planning limit and 1.11 p.u. nameplate ampacity, the violation gate applied throughout.

Parameter	Value
<i>Network and study settings</i>	
Feeder (customers / buses)	IEEE European LV (55 / 906)
Voltage base, LV / MV	0.416 / 11 kV (L–L)
Voltage band (EN 50160)	0.90–1.10 p.u.
Line loading: planning limit / ampacity	1.00 / 1.11 p.u.
Transformer short-term limit	1.20 p.u. of rating
Horizon / resolution	24 h / 0.5 h (48 steps)
Tightening factor β / budget $K_{\max}$	0.6 / 6
<i>Per-home device ratings (sampled ranges)</i>	
PV array rating (mean; range)	7.6 (4.65–10.08) kWp
PV AC peak output (mean; range)	4.8 (3.1–6.6) kW
PV daily yield (mean)	41.5 kWh/home (5.5 kWh/kWp)
Battery capacity / power	5.0–13.5 kWh / 2.5–5.0 kW
EV capacity / charge rate	19–52 kWh / 3.6 kW
EWB rating / tank size	3.0–3.6 kW / 100–150 L
HP rated power	2.2–3.1 kW
Washer-dryer / dishwasher (power, cycle)	2.0 kW, 1.5 h / 1.2 kW, 1.0 h
Refrigerator rated power	0.15 kW

Table 3. Objective weights and scalar settings of the three control layers, applied at every penetration level. A single entry is shared between the coordinated cases; where the two differ, both values are given, and a dash marks a setting that does not apply (Case 2 has no feasibility floors, payback LP, tracking term or settlement; Case 3 has no envelope slack).

Symbol	Role	Layer	Case 2	Case 3
PF	Assumed load power factor (unity)	2, 3	1.0	
ψ^{flex}	Reactive-to-active ratio of the request	3	0.60	0 (active only)
$\lambda_{h,t}^{\text{inc}}$	Compliance rate on imported energy [£/kWh]	1	sigmoid, -0.06 to 0.02	0
w^{trk}	Import-target tracking weight [£/kWh]	1	—	0.05
w^{env}	Envelope-slack penalty [£/kWh]	1	10^3	— (slack fixed at 0)
ε	Margin added to the import cap	2	—	0
w^{need}	Tier-one payback priority	2	—	10
w^{extra}	Tier-two payback priority	2	—	0
γ	Fairness strength in ω_h	2	—	0
η^{rec}	Recoverable fraction of curtailed energy	2	—	1.0
w^{pk}	Feeder-peak penalty in the payback LP	2	—	0.5
w_F	AC-OPF flexibility-deployment weight	3	1.0	5.0
w_L	AC-OPF network-loss weight	3		1.0
w_S	AC-OPF violation-slack weight	3		10^5
κ	Ampacity fraction the request is sized against	3	0.90	1.00
β	Geometric envelope-tightening factor	3		0.60
$K_{\max}$	Maximum dispatch–replay passes	3		6
$\tilde{S}_{\text{tx}}$	Substation transformer rating	3		800 kVA
λ^{fee}	Service fee on import reduction (capped at the allocated cut)	2	—	0.02 £/kWh

which the larger is the early-evening reverse-power peak, which loads the trunk and raises bus voltage simultaneously. This event most clearly distinguishes an import-only scheme from one that also enforces the export limit.

The PV profiles are generated from a clear-sky irradiance model for the feeder's latitude (53.4°N) at the summer solstice, using nine azimuth–tilt combinations (peak plane-of-array irradiance of 855 W/m²), a nominal-operating-cell-temperature (NOCT) cell-temperature model with a $-0.4\%/^{\circ}\text{C}$ power coefficient, 10% system losses and a 96%-efficient inverter clipped at a DC:AC ratio of 1.10, with a clear-sky atmospheric transmittance of 0.80. The array ratings are intentionally large: a pool of typical domestic systems (3.0–6.5 kWp, mean 4.9 kWp) is scaled by 1.55, giving 4.65–10.08 kWp (mean 7.6 kWp); no home reaches the inverter clip on the study day. This stress scenario ensures that the sweep reaches the feeder's physical limit; its implications are discussed in Section 3.14.

Retail prices are taken without modification from a UK supplier. DER-owning homes are placed on Octopus Flux, a solar-and-battery tariff with import prices of £0.1517/kWh off-peak (02:00–05:00), £0.3538/kWh peak (16:00–19:00) and £0.2527/kWh otherwise, and export prices of £0.0466, £0.2925 and £0.1033/kWh in the same bands. Homes without DER are placed on the standard variable tariff at a flat £0.2613/kWh with no export payment. All rates are the region-G (North Western England) values published on 19 September 2026, with import inclusive of VAT; standing charges are omitted. A seeded uniform $\pm£0.006/\text{kWh}$ perturbation is added to the import price of every home at every step to break price ties. The decisive feature is the £0.2925/kWh export band, nearly twice the off-peak import price and more than six times the off-peak export price, which gives a bill-minimising battery a clear incentive to be fully charged at 16:00 and to discharge into the network.

All MILPs are solved with CPLEX 22.1 to a relative optimality gap of 2%, and the per-home problems are solved concurrently. Cost differences below this tolerance are reported as indicative only; this applies to the gross premium at 0% and 25% (Section 3.9) and to the gross-bill comparison in the payback ablation (Section 3.11). All quantities on which the conclusions rest lie well outside this tolerance. The AC-OPF is solved with IPOPT 3.14 and the MUMPS linear solver. All models are implemented in Python 3.11 with Pyomo 6.10, and power flows are solved in OpenDSS through OpenDSSDirect.py 0.9.4.

The objective weights and scalar settings are listed in Table 3. Case 2 applies a sigmoid compliance rate of $-\text{£}0.06$ to $+\text{£}0.02/\text{kWh}$ on imported energy and enforces its import limit through a penalised slack, with no tracking term. Case 3 compensates households entirely through the ex-post settlement (29) with a service fee $\lambda^{\text{fee}} = \text{£}0.02/\text{kWh}$, fixes the envelope slack at zero, and uses a small tracking weight that serves only to break ties among compliant schedules. On the network side, Case 2 couples a reactive component to each request ($Q/P = 0.60$), whereas the proposed request is purely active, and the proposed AC-OPF weights deployed flexibility five times more heavily ($w_F = 5$), which makes its request minimal rather than merely sufficient. Consistent with principle (P3), compliance in Case 3 rests on the hard caps alone, whose feasibility (28) guarantees.

Table 4. Per-home heterogeneity model. All random draws use a fixed seed (42).

Device / quantity	Distribution	Spread
EV arrival / departure time	Fixed (fleet-wide)	18:30 / 08:00
EV initial / departure SoC	Fixed (fleet-wide)	0.20 / 0.80
Battery initial SoC	Fixed (fleet-wide)	0.50
PV array rating	Discrete mix ($\times 1.55$)	4.65–10.08 kWp (Table 2)
PV orientation / tilt	Discrete mix	9 azimuth–tilt combinations
HP comfort band	Uniform	floor 19.5–20.5°C, width 1.8–2.4°C
HP rating / COP	Uniform	2.2–3.1 kW / 2.9–3.4
HP thermal $R^{\text{HP}}, C^{\text{HP}}$	Uniform	2.2–3.0°C/kW, 8.5–11.5 kWh/°C
EWB thermal $R^{\text{EWB}}, C^{\text{EWB}}$	Set by tank size	2.03–3.04°C/kW, 98 h
Hot-water draws	13 archetype profiles	scaled by tank size, ± 90 min shift
EWB tank size / rating	Discrete mix	{100, 120, 150} L / {3.0, 3.3, 3.6} kW
Battery capacity / rate	Discrete mix	{5.0, 6.5, 8.0, 10, 13.5} kWh / C-scaled
Base load	Measured IEEE EULV profiles	55 distinct, 30-min averages
Outdoor temperature	Fixed (fleet-wide)	18–28°C
Wet-appliance requirement	Bernoulli	$p=0.5$ (W/D), 0.6 (DW)
Refrigerator initial temperature	Uniform	$\pm 1.0^{\circ}\text{C}$
DER ownership	Exact-quota bundle	round(55v/100) homes

3.2 Per-home heterogeneity

To avoid the unrealistic synchronism of identical households, the homes differ in demand, device parameters and DER sizes, with every random draw taken from a fixed seed (42). The inflexible base load of each home is one of the 55 distinct measured




Figure 2. Hosting capacity of the three control cases. (a) Summed constraint-violation severity versus DER penetration; the star marks the onset of violation for each case, and the shaded region is the violation-free zone. (b) Resulting hosting capacity, labelled with the limit that sets it. The decentralised baseline is voltage-limited (ceiling 25%), the conventional hierarchical scheme is thermally limited (ceiling 50%), and the proposed framework sustains 100%, the highest level tested.

profiles supplied with the IEEE European LV Test Feeder, averaged to 30 min, scaled by 0.6 and reduced by a further 30% to avoid double-counting the explicitly modelled appliances. Hot-water draws follow 13 archetype profiles, scaled with the tank volume and shifted in time by up to ± 90 min. PV rating and orientation, EWH size and battery size are discrete mixes; the dwelling thermal parameters, HP rating and COP and the comfort bands are uniform draws; the wet-appliance requirement is a Bernoulli trial; and DER ownership is an exact-quota bundle, nested across penetration levels. The outdoor temperature profile, the initial states of charge (0.50 for the battery, 0.20 for the vehicle, departure 0.80) and the vehicle plug-in session (18:30 to 08:00) are deliberately held common across the fleet, so that storage-driven results cannot be attributed to a favourable draw. Table 4 summarises the model.

3.3 Hosting capacity

Fig. 2 presents the central result, and Fig. 3 shows the voltage and loading indicators from which it is derived. On this feeder both network constraints bind, and both are driven by the same event. Under the export-paying tariff, the least-cost schedule charges vehicles and batteries in the 02:00–05:00 low-price window, tops the batteries up from midday PV, and discharges the batteries into the grid when the export price rises at 16:00. Because this behaviour is synchronised across the fleet, the uncoordinated feeder experiences two stress events: coincident early-morning charging and coincident discharge into the export band. At 100% DER, the feeder draws 196 kW at its import peak but exports 225 kW at 16:30. Both events individually overload the trunk (the charging window reaches 1.64 p.u.), but the reverse current determines the daily extremes of 1.73 p.u. loading and 1.117 p.u. voltage, since over-voltage on this feeder is driven by injection. The voltage and thermal limits are therefore reached almost simultaneously, and the binding event occurs in the export direction, which an import-only limit cannot address.

The decentralised baseline is voltage-limited. Its worst bus voltage rises from 1.057 p.u. at 0% to 1.105 p.u. at 50%, breaching the EN 50160 ceiling at 16:00 while line loading remains within ampacity (1.06 p.u.); its hosting ceiling is therefore 25%. At 100%, loading reaches 1.73 p.u., or 156% of nameplate ampacity.

The conventional hierarchical scheme is thermally limited, with a ceiling of 50%. It is identical to the baseline up to 25%, since its envelope engages only after a violation appears. At 50% it clears the violation, reducing the worst bus voltage to 1.084 p.u. and peak loading to 0.90 p.u. At 75% it fails in a manner consistent with its structure: it reduces the import peak from 149 to 91 kW, but the reverse flow at 18:30 increases to 162 kW, because the energy deferred from the import hours remains stored when the export window opens. Peak loading reaches 1.31 p.u. at 75% and 1.68 p.u. at 100%, only marginally below the uncoordinated 1.73 p.u.

The proposed framework is violation-free on both axes at every level tested, with peak loading of 0.70–1.03 p.u. and worst




Figure 3. Network-security indicators versus DER penetration: (a) maximum bus voltage and (b) peak line loading. Dashed lines mark the EN 50160 and thermal limits, and the dotted line in (b) the 90%-of-ampacity reporting base. Stars denote the first penetration level at which each case violates the limit shown in that panel. Only the proposed case remains within both limits at every level. Minimum voltage (≥ 0.981 p.u.) and transformer loading (≤ 0.30 p.u.) are never binding and are omitted.

bus voltage of 1.059–1.093 p.u. It is the only controller to reach 100% penetration, a gain of 75 percentage points over the decentralised ceiling and twice the conventional one at the 25% resolution of the sweep. Two features of this result are notable. First, peak loading rises only from 1.024 to 1.034 p.u. between 75% and 100%, and the daily maximum moves from the tariff window to the overnight charging window (04:30 and 03:00). Second, the result is not obtained by suppressing DER: the fleet still exports 1040 kWh/day, 82% of the uncoordinated total, but spread across the solar day rather than concentrated in three hours.

3.4 Voltage headroom

Fig. 4a shows the full daily voltage band for every case and penetration level. At 50%, uncoordinated export raises the worst bus to 1.1052 p.u. at 16:00, the violation that limits Case 1 to 25%, whereas both coordinated cases hold 1.084 p.u. From 75%, the uncoordinated maximum reaches 1.117 p.u. For the uncoordinated case at 50% and above, the worst voltage occurs within the peak-export tariff window, confirming that the over-voltage and the thermal overload of Section 3.5 are manifestations of the same event. The conventional scheme remains within the band up to 75% but reaches 1.1009 p.u. at 100% under the strict test of Section 2.6; this does not alter its hosting ceiling, which is set earlier by thermal overload, but it indicates that no voltage margin remains. The proposed framework keeps the worst bus at or below 1.093 p.u. and the minimum at or above 1.006 p.u. throughout the sweep. The vertical axis of Fig. 4a is truncated because the lowest voltage in the sweep (0.981 p.u.) is far above the 0.90 p.u. floor; each panel reports its minimum numerically.

3.5 Thermal impact

Fig. 4b shows the daily line-loading band for every case and penetration level. At 0% and 25% the three cases are indistinguishable, with the daily maximum (0.70 and 0.77 p.u.) in the early-evening demand peak. From 50%, the uncoordinated maximum moves into the 16:00–16:30 export window (1.06, 1.46 and 1.73 p.u. at 50%, 75% and 100%) and is set by reverse current. The trunk is overloaded in both directions: the overnight charging window alone reaches 1.31 p.u. at 75% and 1.64 p.u. at 100%.

The import cap of Case 2 acts directly on the overnight peak and removes it, reducing the 00:00–06:00 maximum to 0.90 p.u. and 0.98 p.u. at 75% and 100%. It has no corresponding instrument for the export peak, where its residual violations occur (18:30 at 75% and 16:30 at 100%). Moreover, the deferred energy increases Case 2's evening reverse flow above that of the baseline at 75% (162 against 150 kW). Its peak loading (1.31 and 1.68 p.u.) is therefore only 10% and 3% below the uncoordinated case, and the gap narrows as the import-only mechanism is exhausted. The binding constraint for an import-only scheme is thus the power-flow direction it cannot bound, not the volume of flexibility it procures.

The proposed scheme attains the lowest peak loading from 75% onward: 1.02 and 1.03 p.u. at 75% and 100%, reductions of 30% and 40% relative to the baseline and 22% and 38% relative to Case 2. At 50%, Case 2 is marginally lower (0.90 against




Figure 4. Daily (a) bus-voltage and (b) line-loading bands by control case (columns) and DER penetration (rows): full range (light), 10–90th percentiles (dark) and median (line); the vertical strip is the 16:00–19:00 export window. Dashed lines mark the 1.10p.u. voltage ceiling and the 1.11 p.u. ampacity limit (dotted: 1.00p.u. planning limit); the tinted area above them is the violation region, and violating maxima are printed in red.




Figure 5. Spatial distribution of peak line loading at 100% DER (★ = substation transformer), where 1.00 p.u. is the planning limit and 1.11 p.u. is nameplate ampacity. Overloaded sections are drawn thicker; their number is given in each panel heading: 61 under the decentralised baseline, 57 under the conventional scheme and none under the proposed framework.

0.95 p.u.) because it suppresses import more than required, but both are well within ampacity. At 100%, the aggregator issues export-reduction requests in 14 of the 48 slots (143 kWh/day) and import-reduction requests in 13 slots (171 kWh/day), with the payback returning energy across 22 slots. The export requests fall in two groups: nine in the 08:00–13:00 solar window and five in the 16:00–18:30 tariff window. Acting on both is essential, since the physical limit is violated by reverse power regardless of its cause, and an envelope keyed to the tariff window alone would miss the solar group. The reverse-flow extreme falls from 225 to 135 kW, and the daily maximum moves out of the tariff window to the 03:00 charging flow. The proposed scheme achieves this with less import-side flexibility than Case 2 (Section 3.8), because the AC-OPF procures the minimum flexibility required and the tightening loop applies it only in the violated slots.

Fig. 5 maps peak line loading across the feeder. At 100% DER, 61 line sections exceed nameplate ampacity under the decentralised baseline and 57 under the conventional scheme, against none under the proposed framework. Since a long section of the trunk is overloaded simultaneously, the problem cannot be resolved by reinforcing a single asset. The substation transformer is never binding: its peak loading is 0.30 p.u. of the 800 kVA rating under the uncoordinated baseline at 100% and 0.19 p.u. under the proposed case, far below the 1.2 p.u. short-term threshold. Consequently, the $\sqrt{3}$ looseness of the transformer constraint (Section 2.5) does not affect the results, and a detailed IEEE C57.91 loss-of-life assessment is unnecessary.

3.6 Feeder export and network losses

Fig. 6 shows how the feeder exchanges power with the upstream grid. Under the decentralised baseline, daily household export grows with penetration to 1273 kWh/day at 100%. The conventional scheme leaves it essentially unchanged (1251 kWh/day), since its envelope issues no export cap. The proposed framework tracks the baseline at low penetration and reduces export by 6.9% and 18.3% at 75% and 100% (to 911 and 1040 kWh/day). This reduction is a consequence of bounding the coincident reverse power in the 16:00–19:00 window rather than an objective in itself. PV curtailment is zero in every case and at every penetration level, because the envelopes are anchored on the energy-minimal dispatch (2), which penalises curtailment together with import. The reduced export therefore represents storage-mediated self-supply and deferred discharge rather than wasted generation, as quantified in Section 3.7.

The redistribution also reduces network losses. Under the proposed framework, daily losses are 7.5, 9.0, 16.9, 27.3 and 36.1 kWh/day across the five levels, compared with 7.5, 9.1, 18.1, 33.5 and 54.1 kWh/day for the decentralised baseline, a reduction of 33.3% at 100%. Because losses scale with the square of current and the framework removes the largest currents, reducing the reverse-flow extreme from 225 to 135 kW eliminates roughly two-thirds of the I^2R loss in that window, far more than the off-peak redistribution adds. Case 2 incurs slightly lower losses at 50% (16.0 kWh/day), where it suppresses import more than necessary, but higher losses from 75% onward (29.0 and 44.6 kWh/day), where its unbounded export peak dominates.

The two feeder-head extremes behave differently and are therefore plotted on a single signed axis in Fig. 6(b). At 100%, the

Table 5. Daily fleet energy balance at 100% DER (kWh/day, 55 homes). Inputs are available PV and grid import; outputs are household export, end-use consumption (base load plus the flexible devices), and the net energy absorbed by the EV and battery stores over the day. The final row is the closure residual. Network losses sit outside the household balance and are listed separately for reference.

	Decentralised	Proposed	Difference
<i>In:</i> PV available	2285.0	2285.0	0.0
<i>In:</i> grid import	1131.2	987.8	−143.4
<i>Out:</i> household export	1272.9	1040.4	−232.5
<i>Out:</i> end-use consumption	831.8	906.0	+74.2
<i>Out:</i> EV net absorbed	1266.3	1271.6	+5.3
<i>Out:</i> battery net absorbed	45.3	54.8	+9.5
<i>Out:</i> PV curtailed	0.0	0.0	0.0
Closure residual	0.0	0.0	—
(Network losses, outside the above)	54.1	36.1	−18.0

proposed framework reduces the import peak from 196 to 134 kW, while Case 2 reduces it further, to 108 kW, by suppressing import fleet-wide. For reverse flow the ordering is inverted: 225 kW uncoordinated, 213 kW under Case 2 and 135 kW under the proposed framework, the only scheme that bounds it. A reduction in import peak is therefore not informative for a feeder whose limit is reached on reverse current; at 50% and 75%, Case 2 lowers the import peak while its reverse peak exceeds that of the baseline (84 against 81 kW and 162 against 150 kW). On the binding quantity, the proposed framework reduces the reverse extreme by 40% and the worst-case conductor loading by 40% (1.73 to 1.03 p.u.). These aggregate figures result from a reshaping of the fleet’s daily net power profile: at the stressed levels, the proposed case flattens the early-morning import peak and raises the 16:00–19:00 export trough toward zero by holding stored energy through the tariff window and releasing it when the conductors have capacity. The slot-by-slot form of this reshaping is examined in Section 3.8.

3.7 Where the forgone export goes: a closed energy balance

A reduction in export may correspond to energy consumed on site, stored and released later, curtailed or lost. Table 5 closes the daily energy balance term by term at 100% DER, where the framework intervenes most heavily; the balance closes within the reporting precision at every penetration level.

Three observations follow: First, no generation is wasted: PV curtailment is zero in both cases. Second, the largest destination of the 232.5 kWh/day of forgone export is avoided import: grid import falls by 143.4 kWh/day, or 62% of the export reduction, reflecting energy held in storage through the export window and subsequently used on site. Third, the remaining 89.0 kWh/day (38%) is consumed rather than deferred, comprising 74.2 kWh of additional end-use energy and 14.8 kWh of additional storage losses. About 90% of the additional end-use energy is in the HP; a further 6.6 kWh comes from four optional appliance cycles (one washer-dryer and three dishwasher cycles that the households were not required to run that day), started between 09:00 and 14:30 while their export was at the cap. When the export cap binds and storage is full, surplus PV must either be curtailed or absorbed on site, in a thermal load within its comfort band or an optional appliance cycle; the framework adopts the latter, which is why curtailment is zero, but most of the absorbed energy is dissipated to ambient within the horizon. This amounts to 1.3 kWh per home per day (8% of the proposed case’s daily end-use consumption of 906 kWh) and explains the non-zero gross bill premium. Absorption is preferred because pre-cooling provides some end-use value whereas curtailment provides none, but it is a design choice with a measurable cost. At the system level, this household-side absorption is partly offset by the simultaneous 18.0 kWh/day reduction in network losses.

3.8 Envelope mechanism and deliverability

The two schemes differ in the *shape* of the envelopes they issue (Fig. 7). Case 2 enforces an import limit and leaves export unconstrained; since the violation occurs in the export direction, the scheme cannot act on it and, moreover, defers energy into it. The proposed envelope also enforces the export limit (green hatch). The export band is generous but finite: at 50% DER it integrates to 609 kWh/day against a realised export of 607 kWh/day, so the export limit becomes active only where the network requires it. Together with the import ceiling tightened at the violated slots, this allows the aggregator to keep every conductor and bus within limits while remaining bill-neutral in aggregate.

Fig. 8 compares, at every penetration level, the requested flexibility with that delivered by the HEMS re-dispatch. Procurement is demand-driven: at 0% and 25% the requested volume is negligible and Case 2 reverts to the baseline. At 50% the proposed scheme clears the voltage excursion with 1.5 kWh/day of import reduction and 2.6 kWh/day of export reduction,




Figure 6. Operational energy and demand signature versus DER penetration: (a) daily household export; (b) feeder-head import peak (above zero) and reverse-flow peak (below zero) on one signed axis; (c) network losses as a share of daily import. The conventional scheme attains the lowest import peak (108 kW at 100%) while its reverse peak remains at 213 kW; only the proposed framework bounds the reverse extreme (135 kW). Both coordinated schemes keep the loss share at or below that of the baseline at every penetration level, and the proposed framework attains the lowest share at 100%.

whereas Case 2 requests 18 kWh/day of import reduction. At 75% and 100% the proposed fleet delivers 92.9% and 98.3% of the requested import reduction (60.4 of 65.0 and 168.2 of 171.1 kWh), comparable to Case 2's 92.6% and 96.9%. At 75%, the requested volume includes the verification loop's tightening of the 04:00 import caps (Section 2.5.2).

On the import side, the only side Case 2 can act on, the proposed scheme requests slightly less at 75% (65 against 66 kWh) and 14% less at 100% (171 against 199 kWh), while additionally issuing 38 and 143 kWh/day of export reduction and attaining a 38% lower peak loading. Its total import and export request at 100% (315 kWh/day) exceeds that of Case 2; the advantage therefore lies not in moving less energy but in moving it in the direction in which the violation occurs.


Through the middle of the day, the net target of Case 2 in Fig. 8 remains near zero instead of tracking the exporting baseline. This results from the mapping (41), which acts on import only: an export-reduction request is converted into a raised import target and ceiling. At 10:00 at 100% DER, for example, Case 2's aggregate import target is 146 kW (its aggregate import ceiling 395 kW) while the uncoordinated fleet exports 146 kW, so the net target plotted in Fig. 8 is close to zero. The instruction has no effect, since the fleet is far below the ceiling; the aggregator identifies the export-side problem correctly but lacks a variable with which to address it.

Fulfilment and envelope compliance should be distinguished. Fulfilment measures the fraction of the requested volume that is delivered, whereas envelope compliance measures whether the realised net power respects the hard caps (28). Compliance is 100% at every level, with zero slack in all 55×48 home-time steps, as guaranteed by the feasibility floors; a fulfilment below 100% is therefore a shortfall in voluntary recovery, not a violation of a bound.

The tightening loop (38) is triggered at all three stressed levels and converges in a single pass each time. It acts on one violated *export* slot at 50% (an over-voltage of 1.102 p.u. at 16:00; 28 of 55 export caps are tightened and the others are already at their floors), one *import* slot at 75% (04:00; 32 of 55 import caps, the other 23 being at, or within $2 \mu W$ of, their floors) and one *export* slot at 100% (13:00; 36 of 55 export caps). The tightening applies to the whole feeder but only at the violated step. That the residual violation lies on the export side at two of the three levels directly supports enforcing the export limit, since an import-only loop would have no bound with which to clear it.

3.9 Customer economics

Network compliance carries a tangible cost under an export-paying tariff (Fig. 9). Under the decentralised baseline, the mean gross daily cost falls from £3.81 at 0% to £0.04 at 100%, as 16:00–19:00 export revenue almost entirely offsets household energy costs. Because this revenue corresponds to export that the network cannot absorb, a compliant dispatch must forgo part of it. The proposed framework's gross premium over the baseline is below the solver tolerance at 0% and 25%, and amounts to £0.05, £0.22 and £0.74/day at 50%, 75% and 100%, respectively. This premium reflects forgone export revenue rather than additional purchases, since the proposed case imports 13% less energy than the baseline at 100% (988 against 1131 kWh/day).



both panels at 50% DER penetration

Figure 7. Flexibility-envelope mechanism at 50% DER. (a) Aggregate import and export flexibility envelope with the uncoordinated HEMS schedule and the envelope-constrained dispatch; (b) import-only limit (Case 2) and envelope with enforced import and export limits (Case 3) over a common baseline; only the latter defines an export band (green hatch). Curves above zero show aggregate household import, $\sum_h P_{h,t}^{imp}$, rather than net feeder power, and the region below zero shows the export envelope. Because importing and exporting homes coexist, the fleet still exports 607 kWh/day at this level; the corresponding net power is shown in Fig. 8. The shaded band marks the peak-tariff window.

The conventional scheme raises gross bills by £0.30, £0.46 and £0.65/day at the same levels; at 75% it is therefore both more expensive than the proposed scheme and non-compliant.

These are gross costs, before the ex-post settlement (29). Each home receives a service fee of £0.02/kWh on its import reduction, capped at its allocated cut, and a bill-protection credit equal to any remaining shortfall against its uncoordinated bill. After settlement, all 55 homes are at or below their uncoordinated bills at every penetration level (Fig. 9(a)). The required transfer grows with the displaced export revenue, reaching £43.85/day across the feeder at 100% (£3.43 in service fees and £40.42 in bill protection), or £0.80 per home per day. Capping the fee at the allocated cut means that a household is not paid for reducing more than it was asked. Because households that gain are not charged, the settled fleet-mean cost is marginally below the baseline at 75% and 100%. The fee covers only 8% of the transfer at 100%, which is intended: the guarantee rests on each household's own counterfactual bill rather than on the adequacy of a posted price. The magnitude of the transfer is set by the scenario rather than by the framework, since the same £0.29/kWh export band that makes the uncoordinated feeder unhostable also makes compliance costly.

The contrast with the conventional scheme is the most pronounced economic result. Case 2 leaves 54 of its 55 homes above their uncoordinated bills at every stressed level, by up to £1.17/day at 100%, even after its flexibility payments, because its incentive is priced within the objective rather than settled against each household's baseline. The required per-home settlement at 100% ranges from £0.001 to £2.21/day, a spread of more than three orders of magnitude that no posted price can track; only an ex-post comparison against a counterfactual can protect every household.

3.10 Comfort preservation

All device comfort constraints (EWH temperature (15), indoor and refrigerator temperatures (20) and EV departure state of charge (13)) are imposed as hard HEMS constraints without slack, and the feasibility floors of (28) guarantee that a compliant schedule always exists. Every re-dispatch instance across the sweep was solved within the hard caps, with no envelope slack, so no comfort constraint was relaxed to accommodate an envelope. The network headroom is therefore obtained from scheduling slack, by shifting storage charging, holding stored energy through the export window and adjusting thermal loads within their permitted bands. Comfort neutrality does not imply loss-free operation: as shown in Section 3.7, part of the forgone export is




Figure 8. Flexibility-procurement fulfilment across the DER penetration sweep for the proposed (green) and conventional hierarchical (blue) cases. Each series is the fleet aggregate net power at the connection points (positive for import): the uncoordinated baseline (dotted), the aggregator target (solid) and the delivered re-dispatch (dashed). The baseline falls to -150 kW at 75% and -225 kW at 100%; the proposed target lies above it in that window because the proposed envelope requests an export reduction. The Case 2 net target remains near zero through the export window, because the import-only mapping (41) converts an export-reduction request into a raised import target. The green band is the proposed case's requested import reduction. Annotated volumes are import-side, allowing a like-for-like comparison.

absorbed in thermal loads and dissipated within the horizon. This is a cost in energy, reflected in the gross bill premium, but not in comfort; indoor temperature remains within its band throughout.

3.11 Ablation: which element delivers the hosting gain

To attribute the outcome to individual elements, the three elements that can be disabled independently (the export cap, the payback and the tightening loop) were switched off in turn at 75% and 100% DER, and the tightening loop also at 50%, the lowest level at which it acts, with all other settings unchanged. The feasibility floors are not ablated, because removing them would render the re-dispatch infeasible by construction. Table 6 reports the results.

The export cap provides the hosting-capacity gain. Releasing it (A1) produces violations on both axes at both levels. Peak loading rises from 1.024 to 1.421 p.u. at 75% and to 1.759 p.u. at 100%, which exceeds the uncoordinated baseline (1.734 p.u.). As with Case 2, the energy deferred from the import hours remains stored when the export window opens, and without an export bound its release cannot be prevented. The AC-OPF still requests export reduction (143 kWh/day at 100%), but the means to enforce it is absent; the tightening loop exhausts its budget and reports the residual violation as the feeder's physical boundary.

The payback serves the customer rather than the network. Removing it (A2) leaves the framework compliant and marginally improves the network indicators (peak loading falls from 1.024 to 0.998 p.u. and losses from 27.3 to 26.3 kWh/day at 75%), because the recovery headroom it grants is additional load. The household position, however, deteriorates: the settlement transfer required to keep every home whole rises by 5.8%, from £13.17 to £13.93/day, whereas the change in gross bill (£1.16 to £1.17/day) lies within the optimality tolerance. The payback is therefore justified by bounded energy recovery (Section 2.4.2) rather than by any network requirement.

The tightening loop converts a marginal violation into compliance. Disabling it (A3) leaves small residual violations at every stressed level: 1.1019 p.u. voltage at 50% and 1.1022 p.u. at 100% against the 1.10 p.u. gate, and 1.129 p.u. loading




Figure 9. Customer economics. (a) Per-home *settled* bill change relative to each home’s uncoordinated bill (one point per home; bars show the fleet mean). Under the proposed settlement no home is above zero at any penetration level, whereas under the conventional scheme 54 of 55 homes are above zero from 50% upward. (b) Fleet-mean *gross* cost before settlement. The per-home spread in (b), $-\pounds 5.4$ to $+\pounds 6.3$ /day at 100%, where every home owns the DER bundle, reflects differences in PV, storage and EV sizes and in household demand rather than differences between controllers.

at 75% against the 1.111 p.u. gate. These margins (0.2%, 0.2% and 1.6%, respectively) reflect the divergence between the AC-OPF model and the OpenDSS replay, which a single tightening pass removes. Because the first replay is non-compliant at every stressed level, including 50% (Table 6), the reported hosting capacity without the loop would fall from $\geq 100\%$ to 25%. A1 and A3 therefore fail for opposite reasons: A1 lacks the control instrument, whereas A3 possesses it but does not verify its setting.

The ablation varies one element at a time at two or three penetration levels on a single feeder and tariff; it therefore establishes the function of each element in this scenario rather than the general magnitude of its contribution. Because A1 isolates the effect of the export limit within the proposed framework, it avoids the confounding effect of comparing two differently tuned schemes.

3.12 Computational cost

The wall-clock time of a complete day-ahead cycle on a laptop with an Apple M1 processor (eight cores) and 8 GB of memory, with the $|\mathcal{H}|$ per-home MILPs solved on six parallel workers, is about 40–65 s for the decentralised case, 160–270 s for the conventional hierarchical scheme and 160–300 s for the proposed framework across all penetration levels and repeated runs (wall-clock times on this machine vary by 20–30% between identical runs). A worst case of about five minutes is well within the time available in a day-ahead process, and the second envelope side adds negligible computational burden.

3.13 Synthesis

The results form a coherent chain. The flexibility envelope (22) exposes each home’s enforceable import- and export-reduction capacity; the AC-OPF (30) sizes the minimum relief; the hard caps (28) translate it into attainable per-home bounds; and the tightening loop (38) removes the residual violations, which occur on the export side at 50% and 100%, where an import-only scheme has no bound to tighten. Consequently, every bus and conductor remains within limits up to 100% penetration.

The ablation of Section 3.11 establishes the role of each element: export-limit enforcement provides the hosting capacity, the verification loop secures it, and the payback protects the customer. The cost of the resulting headroom is the net bill premium arising from forgone peak-window export revenue ($\pounds 0.74$ per home per day at 100%), which the ex-post settlement reimburses to households at a fleet cost of $\pounds 43.9$ /day. Table 7 summarises the comparison.

Table 6. Ablation of the proposed framework at 50% (tightening loop only), 75% and 100% DER. Each arm disables exactly one element; all other settings are identical to the full framework. Peak line loading is per unit of the 90%-of-ampacity reporting base, so the violation gate is 1.11 p.u.; the voltage gate is 1.10 p.u. The decentralised baseline is shown for reference at 75% and 100%. The loop column counts the tightening passes applied. Bold marks a violated limit.

DER	Arm	$V_{\max}$ (p.u.)	Peak load (p.u.)	Losses (kWh/d)	Gross cost (£/home/d)	Loop (passes)	Verdict
50%	Full framework	1.084	0.949	16.9	2.06	1	compliant
	A3 no tightening	1.102	0.957	17.8	2.02	0	V
75%	Full framework	1.093	1.024	27.3	1.16	1	compliant
	A1 no export cap	1.104	1.421	28.6	1.07	5	V + I
	A2 no payback	1.093	0.998	26.3	1.17	1	compliant
	A3 no tightening	1.093	1.129	28.1	1.12	0	I
	(decentralised)	1.117	1.461	33.5	0.94	—	V + I
100%	Full framework	1.092	1.034	36.1	0.78	1	compliant
	A1 no export cap	1.104	1.759	41.3	0.40	5	V + I
	A2 no payback	1.092	1.034	34.3	0.78	1	compliant
	A3 no tightening	1.102	1.034	36.7	0.75	0	V
	(decentralised)	1.117	1.734	54.1	0.04	—	V + I

3.14 Limitations and scope

Several limitations bound these findings. First, the study considers a single deterministic day-ahead horizon on one feeder under a high-PV summer day, with a single seeded draw of household parameters, so intra-day forecast error, seasonal variation and feeder-to-feeder diversity are not modelled. The computational results of Section 3.12 vary DER penetration rather than feeder size; a larger feeder enlarges the AC-OPF itself, which is a distinct scaling question not examined here.

Second, the proposed scheme’s voltage margin at 100% (1.092 p.u. against 1.10 p.u.) is narrower than its thermal margin (1.034 p.u. against 1.11 p.u.), so voltage would bind first under a more demanding PV day. The reported hosting capacity of $\geq 100\%$ is a lower bound, since no failure was observed within the tested range, and all hosting capacities are resolved on a 25% grid, so the ceilings of the benchmark schemes and the ratios derived from them are lower bounds at that resolution. The PV ratings (mean 7.6 kWp) are 1.55 times those of a pool of typical domestic systems, and the EV fleet absorbs about 23 kWh per vehicle per day (state of charge 0.20 at arrival and 0.80 at departure), several times typical UK daily driving energy, which intensifies the overnight charging peak. Both choices are deliberate, so that the sweep reaches the feeder’s physical limit; smaller arrays would raise all three hosting ceilings without altering the mechanism that distinguishes the schemes.

Finally, DER ownership, initial storage states, EV plug-in sessions and the outdoor temperature are homogeneous across participating homes; the reactive axis of the envelope is not used, since loads are modelled at unity power factor and inverter Volt/VAr support is disabled; and the fairness weighting (27) is evaluated only at $\gamma = 0$. Heterogeneous ownership, reactive support, fairness calibration and rolling-horizon operation under forecast error are left to future work.

4 Conclusion

This paper proposed a three-layer hierarchical framework for coordinating residential DERs on LV feeders. A home-level HEMS computes cost-optimal and energy-minimal schedules; an aggregator separates enforceable import- and export-reduction capacity from voluntary recovery headroom and converts the network relief request into hard operating envelopes with enforced import and export limits and energy-bounded payback; and a DSO layer sizes the request with an unbalanced AC-OPF and verifies it through an AC-verified tightening loop. Anchoring every cap on the home’s maximum-self-consumption dispatch guarantees that the envelopes are attainable without relaxation, and an ex-post settlement ensures that no household is worse off for complying.

On the IEEE European LV Test Feeder under a commercial UK export-paying tariff, the binding network event is a coincident reverse-power peak in the 16:00–19:00 export band. The proposed framework is the only scheme tested that sustains 100% DER penetration within both the EN 50160 voltage band and conductor ampacity, whereas the decentralised baseline and a conventional import-only hierarchical scheme reach hosting ceilings of 25% and 50%, respectively. The import-only scheme fails for a structural reason: an import cap does not bound reverse power, and the energy it defers remains stored when the export window opens. The hosting gain is attributable to enforcing the export limit rather than to the volume of procured flexibility. At full penetration, the framework requests 86% of the conventional scheme’s import-side volume, additionally requests 143 kWh/day of export-side reduction, and attains a 38% lower peak conductor loading. The ablation study confirms

Table 7. Performance summary across the DER penetration sweep (per day, 55-customer IEEE European LV feeder). *Hierarchical* denotes the conventional import-only hierarchical scheme (Case 2). Green denotes compliance and bold red a violation of a hard operating limit. Peak line loading is expressed per unit of the 90%-of-ampacity planning limit (nameplate ampacity = 1.11 p.u.), and the voltage test is applied to unrounded values, so 1.1009 p.u. is a violation. Peak reverse flow is the most negative feeder-head net demand; household export is the sum of individual dwelling exports. Gross and settled costs are before and after the ex-post settlement (29), and the settlement transfer is the total budget required. Envelope compliance is the fraction of home-time steps within the hard caps (28). The proposed scheme’s hosting capacity is a lower bound, since 100% is the highest level tested.

Metric	Controller	0%	25%	50%	75%	100%
Peak line loading (p.u.; 1.11 = ampacity)	Decentralised	0.70	0.77	1.06	1.46	1.73
	Hierarchical	0.70	0.77	0.90	1.31	1.68
	Proposed	0.70	0.77	0.95	1.02	1.03
Thermal overload	Decentralised	no	no	no	yes	yes
	Hierarchical	no	no	no	yes	yes
	Proposed	no	no	no	no	no
$V_{\max}$ (p.u.; 1.10 = EN 50160 ceiling)	Decentralised	1.057	1.075	1.105	1.117	1.117
	Hierarchical	1.057	1.075	1.084	1.099	1.101
	Proposed	1.059	1.075	1.084	1.093	1.092
$V_{\min}$ (p.u.; 0.90 = EN 50160 floor)	Decentralised	1.015	1.020	1.004	0.992	0.981
	Hierarchical	1.015	1.020	1.012	1.011	1.001
	Proposed	1.015	1.020	1.006	1.015	1.006
Peak reverse flow (kW, feeder head)	Decentralised	0	20	81	150	225
	Hierarchical	0	20	84	162	213
	Proposed	0	20	80	112	135
Household export (kWh/day)	Decentralised	0	300	628	979	1273
	Hierarchical	0	300	620	965	1251
	Proposed	0	299	607	911	1040
Network losses (kWh/day)	Decentralised	7.5	9.1	18.1	33.5	54.1
	Hierarchical	7.5	9.1	16.0	29.0	44.6
	Proposed	7.5	9.0	16.9	27.3	36.1
Mean gross cost (£/day)	Decentralised	3.81	3.03	2.01	0.94	0.04
	Hierarchical	3.81	3.03	2.31	1.40	0.69
	Proposed	3.82	3.03	2.06	1.16	0.78
Mean settled cost (£/day)	Decentralised	3.81	3.03	2.01	0.94	0.04
	Hierarchical	3.81	3.03	2.35	1.40	0.66
	Proposed	3.81	3.03	2.01	0.92	−0.02
Settlement transfer, proposed (£/day)		0.25	0.36	3.01	13.17	43.85
Hosting capacity	Decentralised 25%, Hierarchical 50%, Proposed $\geq 100\%$					
Proposed envelope compliance (%)		100	100	100	100	100

that removing the export bound alone raises peak loading to 1.76 p.u., above the uncoordinated case, while the tightening loop removes a residual model-to-simulation divergence that would otherwise limit hosting capacity to 25%. Network losses fall by 33% relative to the uncoordinated baseline, PV curtailment is zero, household export decreases by only 18%, and all comfort constraints are satisfied without slack. A complete day-ahead cycle requires at most about five minutes of computation.

Network security under this tariff entails a net bill premium, arising from forgone peak-window export revenue, of £0.74 per home per day at full penetration, which the ex-post settlement reimburses to households, so that the cost is borne by the party procuring the service rather than by the households that provide it. These results indicate that, where tariffs remunerate export in a fixed window, the export limit governs network security, and hierarchical schemes must enforce both the import and export limits of the operating envelope to secure high DER penetration. Future work will examine alternative tariff structures, multi-feeder and medium-voltage coordination, intra-day envelope updates under forecast uncertainty, and scalability to larger networks.

Funding

This work was funded by the Deanship of Graduate Studies and Scientific Research at Qassim University (www.qu.edu.sa), grant QU-APC-2026.

Author contributions

O.A. conceived the framework, developed the models and simulation implementation, analysed and interpreted the results, and wrote and revised the manuscript.

Data availability

The IEEE European Low-Voltage Test Feeder, including its network data and residential load profiles, is publicly available from the IEEE PES Test Feeder repository (<https://cmte.ieee.org/pes-testfeeders/resources/>). The electricity tariff rates used are the publicly published Octopus Energy Flux and standard variable tariff rates for region G (North Western England) as of 19 September 2026. The datasets generated and analysed during the current study, including household schedules, flexibility envelopes, dynamic operating envelopes, network power-flow results and settlement outputs, are not publicly deposited at this time because they were produced with a simulation framework that underpins an ongoing, unpublished extension of this work by the corresponding author. The datasets underlying the specific results reported in this manuscript are available from the corresponding author upon reasonable request.

Code availability

The Python/Pyomo implementation of the three-layer framework, comprising the household HEMS, the aggregator flexibility-coordination and envelope-allocation model, the AC-OPF-based dynamic operating envelope computation and the ex-post settlement, together with the OpenDSS interfacing scripts used to generate the results reported in this manuscript, is not publicly deposited at this time. The codebase underlies an ongoing, unpublished extension of this work by the corresponding author, and public release at this stage would risk premature disclosure of methods and findings from that separate study. The code underlying the specific results reported in this manuscript is available to the editors and reviewers in confidence for the purpose of peer review, and from the corresponding author upon reasonable request, subject to a data-use agreement restricting redistribution.

Competing interests

The author declares no competing interests.

References

1. National Grid Electricity System Operator. Future energy scenarios: ESO pathways to net zero. Tech. Rep., National Grid ESO, Warwick, UK (2024). <https://www.neso.energy/document/321041/download>.
2. Australian Energy Market Operator. 2024 integrated system plan for the national electricity market. Tech. Rep., AEMO, Melbourne (2024). <https://www.aemo.com.au/-/media/files/major-publications/isp/2024/2024-integrated-system-plan-isp.pdf>.

3. Wirth, H. Recent facts about photovoltaics in Germany. Tech. Rep., Fraunhofer Institute for Solar Energy Systems ISE, Freiburg (2026). Version of 20 August 2026. <https://www.ise.fraunhofer.de/content/dam/ise/en/documents/publications/studies/recent-facts-about-photovoltaics-in-germany.pdf>.
4. Ofgem. Smart export guarantee (SEG). Office of Gas and Electricity Markets (2020). Scheme in force from 1 January 2020. <https://www.ofgem.gov.uk/environmental-and-social-schemes/smart-export-guarantee-seg>.
5. Sun, K., Sarker, M. R. & Ortega-Vazquez, M. A. Statistical characterization of electric vehicle charging in different locations of the grid. In *IEEE Power Energy Soc. Gen. Meet.*, DOI: [10.1109/PESGM.2015.7285794](https://doi.org/10.1109/PESGM.2015.7285794) (2015).
6. CENELEC. EN 50160:2022: Voltage characteristics of electricity supplied by public electricity networks. European Committee for Electrotechnical Standardization (2022).
7. Stetz, T., Marten, F. & Braun, M. Improved low voltage grid-integration of photovoltaic systems in Germany. *IEEE Trans. Sustain. Energy* **4**, 534–542 (2013).
8. Navarro-Espinosa, A. & Ochoa, L. F. Probabilistic impact assessment of low carbon technologies in LV distribution systems. *IEEE Trans. Power Syst.* **31**, 2192–2203 (2016).
9. McKenna, K. & Keane, A. Residential Load Modeling of Price-Based Demand Response for Network Impact Studies. *IEEE Trans. Smart Grid* **7**, 2285–2294, DOI: [10.1109/TSG.2015.2437451](https://doi.org/10.1109/TSG.2015.2437451) (2016).
10. Paterakis, N. G., Erdinç, O. & Catalão, J. P. S. An overview of demand response: Key-elements and international experience. *Renew. Sustain. Energy Rev.* **69**, 871–891 (2017).
11. Huang, S., Wu, Q., Cheng, L. & Liu, Z. Optimal reconfiguration-based dynamic tariff for congestion management and line loss reduction in distribution networks. *IEEE Trans. Smart Grid* **7**, 1295–1303 (2016).
12. Hussain, S., Azim, M. I., Lai, C. & Eicker, U. Smart home integration and distribution network optimization through transactive energy framework – a review. *Appl. Energy* **395**, 126193, DOI: [10.1016/j.apenergy.2025.126193](https://doi.org/10.1016/j.apenergy.2025.126193) (2025).
13. Hussain, S., El-Bayeh, C. Z., Lai, C. & Eicker, U. Multi-Level Energy Management Systems Toward a Smarter Grid: A Review. *IEEE Access* **9**, 71994–72016, DOI: [10.1109/ACCESS.2021.3078082](https://doi.org/10.1109/ACCESS.2021.3078082) (2021).
14. Hussain, S., Alrumayh, O. & Lai, C. Integrating device-level flexibility in home energy management systems for prosumers. *IEEE Access* **13**, 157126–157136, DOI: [10.1109/ACCESS.2025.3606499](https://doi.org/10.1109/ACCESS.2025.3606499) (2025).
15. Hilshey, A. D., Rezaei, P., Hines, P. D. & Frolik, J. Electric vehicle charging: Transformer impacts and smart, decentralized solutions. In *2012 IEEE Power and Energy Society General Meeting*, 1–8 (IEEE, 2012).
16. Hussain, S., Lai, C. & Eicker, U. Flexibility: Literature review on concepts, modeling, and provision method in smart grid. *Sustain. Energy, Grids Networks* **35**, 101113, DOI: [10.1016/j.segan.2023.101113](https://doi.org/10.1016/j.segan.2023.101113) (2023).
17. Bollen, M. H. J. & Hassan, F. *Integration of Distributed Generation in the Power System* (Wiley-IEEE Press, 2011).
18. Ruiz-Rodriguez, F. J., Hernandez, J. C. & Jurado, F. Probabilistic load flow for photovoltaic distributed generation using the Cornish–Fisher expansion. *Electr. Power Syst. Res.* **89**, 129–138 (2012).
19. Yan, R. & Saha, T. K. Investigation of voltage stability for residential customers due to high photovoltaic penetrations. *IEEE Trans. Power Syst.* **27**, 651–662 (2012).
20. Ismael, S. M., Abdel Aleem, S. H. E., Abdelaziz, A. Y. & Zobaa, A. F. State-of-the-art of hosting capacity in modern power systems with distributed generation. *Renew. Energy* **130**, 1002–1020 (2019).
21. Baran, M. E. & Wu, F. F. Network reconfiguration in distribution systems for loss reduction and load balancing. *IEEE Trans. Power Deliv.* **4**, 1401–1407 (1989).
22. Turitsyn, K., Šulc, P., Backhaus, S. & Chertkov, M. Options for control of reactive power by distributed photovoltaic generators. *Proc. IEEE* **99**, 1063–1073 (2011).
23. Fortenbacher, P., Ulbig, A. & Andersson, G. Optimal placement and sizing of distributed battery storage in low voltage grids using receding horizon control strategies. *IEEE Trans. Power Syst.* **33**, 2383–2394 (2018).

24. Australian Energy Market Commission. Access, pricing and incentive arrangements for distributed energy resources: Rule determination. Tech. Rep., AEMC, Sydney (2021). <https://www.aemc.gov.au/rule-changes/access-pricing-and-incentive-arrangements-distributed-energy-resources>.
25. Petrou, K. *et al.* Ensuring distribution network integrity using dynamic operating limits for prosumers. *IEEE Trans. Smart Grid* **12**, 3877–3888 (2021).
26. Liu, M. Z., Ochoa, L. F., Wong, P. K. C. & Theunissen, J. Using OPF-based operating envelopes to facilitate residential DER services. *IEEE Trans. Smart Grid* **13**, 4494–4504, DOI: [10.1109/TSG.2022.3188927](https://doi.org/10.1109/TSG.2022.3188927) (2022).
27. Attarha, A., Scott, P. & Thiébaux, S. Affinely adjustable robust ADMM for residential DER coordination in distribution networks. *IEEE Trans. Smart Grid* **11**, 1620–1629, DOI: [10.1109/TSG.2019.2941235](https://doi.org/10.1109/TSG.2019.2941235) (2020).
28. Christakou, K., Le Boudec, J.-Y., Paolone, M. & Tomozei, D.-C. Efficient computation of sensitivity coefficients of node voltages and line currents in unbalanced radial electrical distribution networks. *IEEE Trans. Smart Grid* **4**, 741–750 (2013).
29. Dobbe, R. *et al.* Toward distributed energy services: Decentralizing optimal power flow with machine learning. *IEEE Trans. Smart Grid* **11**, 1296–1306 (2020).
30. Wickramasinghe, A., Vilathgamuwa, M., Nourbakhsh, G. & Corry, P. A review of dynamic operating envelopes: Computation, applications and challenges. *Modelling* **6**, 29, DOI: [10.3390/modelling6020029](https://doi.org/10.3390/modelling6020029) (2025).
31. Barzegar, A. *et al.* Dynamic operating envelopes in distribution grids: Current advances, challenges, and future opportunities. *Appl. Energy* **421**, 128172, DOI: [10.1016/j.apenergy.2026.128172](https://doi.org/10.1016/j.apenergy.2026.128172) (2026).
32. Distributed Energy Integration Program. Navigating dynamic operating envelopes: Establishing a shared language. Tech. Rep., Distributed Energy Integration Program (DEIP), Australian Renewable Energy Agency (2026). <https://arena.gov.au/assets/2026/01/Final-DEIP-DOEs-Report-Feb-2026.pdf>.
33. Lankeshwara, G., Sharma, R., Alam, M. R., Yan, R. & Saha, T. K. Development and validation of a dynamic operating envelopes-enabled demand response scheme in low-voltage distribution networks. *Appl. Energy* **382**, 125150, DOI: [10.1016/j.apenergy.2024.125150](https://doi.org/10.1016/j.apenergy.2024.125150) (2025).
34. Liu, Z. *et al.* Two-stage optimal scheduling of electric vehicle charging based on transactive control. *IEEE Trans. Smart Grid* **10**, 2948–2958, DOI: [10.1109/TSG.2018.2815593](https://doi.org/10.1109/TSG.2018.2815593) (2019).
35. Li, S., Zhang, D., Roget, A. B. & O'Neill, Z. Integrating home energy simulation and dynamic electricity price for demand response study. *IEEE Trans. Smart Grid* **5**, 779–788, DOI: [10.1109/TSG.2013.2279110](https://doi.org/10.1109/TSG.2013.2279110) (2014).
36. Molzahn, D. K. *et al.* A survey of distributed optimization and control algorithms for electric power systems. *IEEE Trans. Smart Grid* **8**, 2941–2962, DOI: [10.1109/TSG.2017.2720471](https://doi.org/10.1109/TSG.2017.2720471) (2017).
37. Nguyen, D. T. & Le, L. B. Optimal bidding strategy for microgrids considering renewable energy and building thermal dynamics. *IEEE Trans. Smart Grid* **5**, 1608–1620 (2014).
38. Aussel, D., Brotcorne, L., Lepaul, S. & von Niederhäusern, L. A trilevel model for best response in energy demand-side management. *Eur. J. Oper. Res.* **281**, 299–315 (2020).
39. Li, L. Coordination between smart distribution networks and multi-microgrids considering demand side management: A trilevel framework. *Omega* **102**, 102326 (2021).
40. Wang, H. *et al.* Distributed coordinative transaction of a community integrated energy system based on a tri-level game model. *Appl. Energy* **295**, 116972 (2021).
41. Almassalkhi, M. *et al.* Hierarchical, grid-aware, and economically optimal coordination of distributed energy resources in realistic distribution systems. *Energies* **13**, 6399, DOI: [10.3390/en13236399](https://doi.org/10.3390/en13236399) (2020).
42. Mak, D. & Choi, D.-H. Optimization framework for coordinated operation of home energy management system and volt-var optimization in unbalanced active distribution networks considering uncertainties. *Appl. Energy* **276**, 115495 (2020).

43. Hussain, S., El-Bayeh, C. Z., Menon, R. P., Lai, C. & Eicker, U. Flexibility based coordination framework for three-level energy management system. In *2022 19th International Bhurban Conference on Applied Sciences and Technology (IBCAST)*, 568–572 (IEEE, 2022).
44. Alrumayh, O. Decentralized, centralized, and hierarchical coordination of residential DERs in three-phase unbalanced distribution networks: A comparative analysis. *Sci. Reports* DOI: [10.1038/s41598-026-64070-1](https://doi.org/10.1038/s41598-026-64070-1) (2026).
45. Hussain, S., Alrumayh, O., Parakkal Menon, R., Lai, C. & Eicker, U. Novel incentive-based multi-level framework for flexibility provision in smart grids. *IEEE Trans. Smart Grid* **15**, 1594–1607, DOI: [10.1109/TSG.2023.3302520](https://doi.org/10.1109/TSG.2023.3302520) (2024).
46. Hussain, S., Azim, M. I., Lai, C. & Eicker, U. Multi-stage optimization for energy management and trading for smart homes considering operational constraints of a distribution network. *Energy Build.* **301**, 113722, DOI: [10.1016/j.enbuild.2023.113722](https://doi.org/10.1016/j.enbuild.2023.113722) (2023).
47. Alrumayh, O. & Bhattacharya, K. Flexibility of Residential Loads for Demand Response Provisions in Smart Grid. *IEEE Trans. Smart Grid* **10**, 6284–6297, DOI: [10.1109/TSG.2019.2901191](https://doi.org/10.1109/TSG.2019.2901191) (2019).
48. Rastegar, M., Fotuhi-Firuzabad, M. & Moeini-Aghaie, M. Developing a two-level framework for residential energy management. *IEEE Trans. Smart Grid* **9**, 1707–1717, DOI: [10.1109/TSG.2016.2598754](https://doi.org/10.1109/TSG.2016.2598754) (2018).
49. Mohsenian-Rad, A.-H., Wong, V. W., Jatskevich, J., Schober, R. & Leon-Garcia, A. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Trans. Smart Grid* **1**, 320–331 (2010).
50. Sortomme, E. & El-Sharkawi, M. A. Optimal charging strategies for unidirectional vehicle-to-grid. *IEEE Trans. Smart Grid* **2**, 131–138 (2011).
51. Nottrott, A., Kleissl, J. & Washom, B. Energy dispatch schedule optimization and cost benefit analysis for grid-connected, photovoltaic-battery storage systems. *Renew. Energy* **55**, 230–240 (2013).
52. Kondoh, J., Lu, N. & Hammerstrom, D. J. An evaluation of the water heater load potential for providing regulation service. *IEEE Trans. Power Syst.* **26**, 1309–1316 (2011).
53. Mathieu, J. L., Dyson, M. & Callaway, D. S. Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations. In *Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings*, 1–189–1–203 (2012).
54. Gottwalt, S., Ketter, W., Block, C., Collins, J. & Weinhardt, C. Demand side management—a simulation of household behavior under variable prices. *Energy Policy* **39**, 8163–8174 (2011).
55. Perfumo, C., Kofman, E., Braslavsky, J. H. & Ward, J. K. Load management: Model-based control of aggregate power for populations of thermostatically controlled loads. *Energy Convers. Manag.* **55**, 36–48, DOI: [10.1016/j.enconman.2011.10.019](https://doi.org/10.1016/j.enconman.2011.10.019) (2012).
56. Aghaei, J. & Alizadeh, M.-I. Demand response in smart electricity grids equipped with renewable energy sources: A review. *Renew. Sustain. Energy Rev.* **18**, 64–72 (2013).
57. Rigoni, V. & Keane, A. Open-DSOPF: An open-source optimal power flow formulation integrated with OpenDSS. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, 1–5, DOI: [10.1109/PESGM41954.2020.9282125](https://doi.org/10.1109/PESGM41954.2020.9282125) (2020).
58. Tesla, Inc. Powerwall 2 AC datasheet. Tech. Rep., Tesla, Inc. (2019). <https://www.tesla.com/powerwall>.
59. U.S. Federal Highway Administration. 2017 national household travel survey. Tech. Rep., U.S. Department of Transportation (2018). <https://nhts.orl.gov>.
60. Du, P. & Lu, N. Appliance commitment for household load scheduling. *IEEE Trans. Smart Grid* **2**, 411–419, DOI: [10.1109/TSG.2011.2140344](https://doi.org/10.1109/TSG.2011.2140344) (2011).
61. Paterakis, N. G., Erdinç, O., Bakirtzis, A. G. & Catalão, J. P. Optimal household appliances scheduling under day-ahead pricing and load-shaping demand response strategies. *IEEE Trans. Ind. Informatics* **11**, 1509–1519, DOI: [10.1109/TII.2015.2438534](https://doi.org/10.1109/TII.2015.2438534) (2015).

62. Fotouhi Ghazvini, M. A., Soares, J., Abrishambaf, O., Castro, R. & Vale, Z. Demand response implementation in smart households. *Energy Build.* **143**, 129–148, DOI: [10.1016/j.enbuild.2017.03.020](https://doi.org/10.1016/j.enbuild.2017.03.020) (2017).
63. Garcia, P. A., Pereira, J. L. R., Carneiro, S., Da Costa, V. M. & Martins, N. Three-phase power flow calculations using the current injection method. *IEEE Trans. Power Syst.* **15**, 508–514 (2000).
64. Wächter, A. & Biegler, L. T. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.* **106**, 25–57 (2006).
65. IEEE PES Distribution System Analysis Subcommittee. European Low Voltage Test Feeder. IEEE PES Test Feeder repository (2015). <https://cmte.ieee.org/pes-testfeeders/resources/>.
66. Dugan, R. C. & McDermott, T. E. An open source platform for collaborating on smart grid research. In *2011 IEEE Power and Energy Society General Meeting*, 1–7, DOI: [10.1109/PES.2011.6039829](https://doi.org/10.1109/PES.2011.6039829) (2011).