

Multi-Level Adaptive Energy Coordination (MAEC)

A Reliability-Constrained, Value-Sharing Architecture for Renewable-Dominant Power Systems, with an Evaluation Design for India

Deepak
Independent researcher

Abstract—Renewable-dominant electricity systems must coordinate a rapidly growing number of small, distributed resources (rooftop solar, batteries, electric vehicles, smart air conditioners, irrigation pumps and thermal loads) with distribution networks and the bulk power system. In India, installed renewable capacity reached about 288.58 GW by 30 June 2026, rooftop solar has been installed on more than five million homes under PM Surya Ghar, and over eleven lakh standalone solar pumps have been deployed under PM-KUSUM. These developments create a coordination problem that spans timescales from sub-second frequency response to multi-year investment, and spatial scales from a single device to the national grid. Existing approaches, including microgrids, demand response, virtual power plants, transactive energy and hierarchical control, each address parts of this problem. This white paper proposes Multi-Level Adaptive Energy Coordination (MAEC), a five-layer architecture (device, building, community hub, distribution network and state, region and nation) connected by shared decision logic. It identifies three candidate mechanisms that may extend existing work: (M1) cross-layer value allocation that pays a resource for each distinct service it provides without double counting; (M2) reliability guardrails exchanged between adjacent layers so that no layer's optimisation violates another layer's security limits; and (M3) adaptive, tiered commitments (firm, likely, possible) that can be revised in an orderly way as forecasts and network conditions change. The paper gives a mathematical formulation with a consequence-weighted cost of unserved energy, a distributed solution approach, an Indian case-study design, two established baselines, eight evaluation metrics, a stress-scenario set and an ablation protocol that attributes any improvement to coordination rather than to additional resources. No empirical results are reported. The contribution of this document is a precisely stated, falsifiable research design and a preliminary novelty assessment intended to guide simulation, pilot testing and peer review.

Index Terms—Distributed energy resources; hierarchical coordination; transactive energy; virtual power plants; demand flexibility; value allocation; Shapley value; reliability; resilience; distribution networks; India.

EXECUTIVE SUMMARY

The problem: India's power system is adding solar capacity faster than it is adding the flexibility needed to use that capacity at the right time and place. Midday solar output can exceed what networks can absorb, while evening peaks, driven increasingly by cooling, stress distribution transformers and raise power purchase costs. Crores of small devices could provide the missing flexibility, but they are not coordinated across the levels of the system, and their owners are rarely paid for the full value they could create.

The proposal: MAEC organises coordination in five layers. Each layer optimises its own objectives within limits ("guardrails") published by its neighbours, makes tiered commitments to the layer above, and is paid through a value-allocation rule that recognises services at more than one layer. The design reuses established methods wherever possible; its potential contribution lies in three coordination rules and in their combination.

What is established and what is not.

Element	Status	What this research must show
Coordinating solar, storage and flexible demand	Established	A specific, measurable improvement
Hierarchical and distributed energy control	Established	A distinct coordination rule, not a new diagram

Virtual power plants and transactive energy	Established	Superior or complementary performance
One framework jointly addressing cross-layer value, guardrails and revisable commitments	Candidate research opportunity	Added value over strong baselines
A specific protocol or control algorithm	Unknown until investigated	Literature, patent search, specification and testing

How it will be tested: On realistic Indian test networks, MAEC will be compared with (A) conventional operation and (B) a strong established method using the same resources, data and reliability requirements, under normal and stress scenarios. Ablation tests will isolate the contribution of each mechanism.

What decision-makers can do now: Low-regret actions include enabling time-of-day tariffs with protection for small consumers, specifying an open signal standard for new air conditioners, chargers, pumps and water heaters, and authorising regulatory sandboxes for transformer-level community hubs.

I. INTRODUCTION AND RESEARCH MOTIVATION

1.1 Context

India's renewable expansion has been rapid. According to information provided by the Ministry of New and Renewable Energy (MNRE) to Parliament, installed renewable capacity reached about 288.58 GW on 30 June 2026, compared with about 76 GW in 2014; solar accounted for about 162 GW, wind about 57 GW, large hydropower about 57 GW and bio-power about 12 GW [1]. About 29 GW of solar and wind capacity was added in the first half of 2026 alone, with rooftop solar the fastest-growing segment [2]. Under PM Surya Ghar: Muft Bijli Yojana, launched in February 2024, rooftop systems had been installed on more than five million households with about 15 GW of cumulative capacity by mid-2026 [3]. Under PM-KUSUM, about 11.69 lakh standalone solar pumps had been installed by July 2026, while decentralised ground-mounted plants under Component A had reached about 2 GW of a 10 GW target [4]. These figures describe capacity. They do not describe whether energy is available at the hour and location

where it is needed. Solar output peaks around midday, while system demand in many Indian states peaks in the evening and, during hot seasons, late at night. The resulting net-load profile, known as the "duck curve" [8], produces a deep midday trough and a steep evening ramp (Figure 1).




Figure 1. Illustrative net-load ("duck") curve. Demand minus solar output produces a midday trough and a steep evening ramp. Values are schematic, not measured data.

1.2 Why distribution-level coordination now matters

Three trends shift the coordination challenge toward the distribution network and behind-the-meter resources:

1. Reverse and variable power flows. Distribution networks designed for one-way flow now host large volumes of rooftop export at midday, which can raise local voltages and trigger inverter disconnection.
2. Concentrated cooling peaks. Air-conditioning adoption causes coincident evening and night loads that overload distribution transformers sized for earlier load patterns.
3. Growth of shiftable loads. Water pumping, water heating, cold storage and electric-vehicle charging are among the fastest-growing loads and are also among the most shiftable, provided that devices can receive and act on coordination signals.

Previous analyses of renewable integration in India, including work by the International Energy Agency [5, 6] and RMI [7], have identified flexibility from storage, demand and buildings as central to cost-effective integration. The question addressed here is how such flexibility can be coordinated across levels, fairly and reliably, at the scale of crores of devices.

1.3 Research question

The central research question is:

Central research question: Can a multi-level coordination method decide, better than established methods, when a distributed resource should serve its owner, support its local network, or provide services to the wider grid, while preserving reliability and reducing total system cost under equivalent resources and constraints?

1.4 Candidate contributions

This paper proposes, but does not yet claim to have demonstrated, the following contributions:

1. An architecture specification (Section 4): five coordination layers with defined decision rights, information interfaces and privacy boundaries.
2. Three coordination mechanisms (Sections 4 and 5): cross-layer value allocation (M1), inter-layer reliability guardrails (M2) and adaptive tiered commitments (M3).
3. A consequence-weighted reliability formulation (Section 5.8) that values unserved energy by load priority rather than by a single average value.
4. An evaluation design (Sections 6 and 7) with baselines, metrics, stress scenarios and ablation tests that can attribute improvements specifically to coordination.
5. A preliminary novelty assessment (Section 3) and a reproducible procedure for completing it.

1.5 Structure of the paper

Section 2 defines the problem and the limitations of current approaches. Section 3 reviews related work and states the novelty hypothesis. Section 4 describes the architecture and mechanisms. Section 5 gives the mathematical formulation. Section 6 describes the Indian case-study systems and data. Section 7 sets out baselines, experiments and metrics. Section 8 states testable hypotheses and illustrative worked examples. Section 9 discusses economic, regulatory and operational feasibility. Section 10 addresses limitations, threats to validity and a roadmap.

II. PROBLEM DEFINITION AND LIMITATIONS OF CURRENT APPROACHES

2.1 The multi-level coordination problem

Let a power system consist of a set of distributed resources embedded in buildings, which are connected

through low-voltage networks to distribution transformers, through medium-voltage feeders to substations, and through transmission to the bulk system. Each resource can, at a given time, provide one or more of the following services: (i) meeting its owner's demand or comfort requirements, (ii) relieving a local network constraint, (iii) relieving a distribution-system constraint or reducing the distribution company's (DISCOM's) procurement cost, and (iv) supporting system-wide balance, reserves or ramping. The coordination problem is to decide, for every resource and time interval, which services to provide, subject to physical limits, owner preferences and the security of every layer, and to remunerate resources in a way that motivates participation without paying twice for the same benefit.

2.2 Four characteristic failure modes

Four failure modes illustrate why single-layer optimisation is insufficient:

Failure mode	Mechanism	Consequence
Local export overvoltage	Many rooftop systems on one feeder export simultaneously at midday	Voltage rise, inverter tripping, lost generation
Herding or rebound peak	Many devices respond identically to the same price signal (for example, a night tariff boundary)	A new peak that overloads transformers
Locally selfish operation	A community battery always reserved for local backup	Idle flexibility during wider system scarcity
Globally selfish operation	Aggregated household batteries fully dispatched for the evening system peak	Owners lose backup when a local outage follows

In each case, a decision that is optimal at one layer imposes a cost or risk on another layer that is neither priced nor constrained.

2.3 Multiple timescales

Coordination must also be consistent across timescales:

Timescale	Coordination task	Typical actor
Sub-second to seconds	Frequency and voltage response	Device controllers, inverters
Seconds to minutes	Balancing, congestion relief	Batteries, fast flexible loads

15 minutes to hours	Intraday scheduling	Building and community energy managers
Day-ahead	Forecasting and commitment	DISCOMs, aggregators, load dispatch centres
Weeks to seasons	Water, fuel and monsoon planning	State and national planners
Years	Network and generation investment	Planners, regulators

Indian scheduling and settlement operate in 15-minute blocks (96 per day), which provides a natural operational interval for the coordination layers above the device level.

2.4 Limitations of current approaches

Approach	Strength	Typical limitation addressed by this research
Centralised optimisation and control	Can reach the optimum for a well-specified model	Communication, computation and privacy burdens at the scale of millions of devices; single point of failure
Price-based and transactive coordination	Decentralised; uses local information	Herding, exposure of consumers to complex prices, weak treatment of local network limits
Virtual power plants	Aggregates small resources into market-sized products	Often optimised for wholesale or system services without explicit local network guardrails
Microgrids	Strong local reliability and islanding	May withhold flexibility from the wider system
Hierarchical or multi-level energy management	Matches decision layers to physical layers	Remuneration across layers and revisable inter-layer commitments are often outside scope

These limitations are stated as hypotheses about the typical treatment in the literature. Section 3 describes how they will be verified paper by paper.

III. RELATED WORK AND NOVELTY STATEMENT

3.1 Microgrids and hierarchical control

Microgrids, introduced as a coordinated cluster of loads and micro-sources able to operate connected to or islanded from the main grid [9], are commonly controlled through hierarchical primary, secondary and tertiary layers [10]. This work establishes that layered control can separate fast local stability from slower economic optimisation. MAEC adopts this principle but extends the hierarchy beyond the microgrid boundary to the DISCOM, state and national layers.

3.2 Demand response and aggregated load control

Demand response has been studied extensively as both incentive-based and price-based programmes [12]. Aggregated control of large populations of thermostatic and other loads can provide fast, accurate services when aggregate dynamics are modelled appropriately [11]. These results support MAEC's assumption that populations of small loads can provide dependable flexibility, provided that baselines, comfort limits and rebound effects are handled.

3.3 Virtual power plants

The virtual power plant concept aggregates distributed resources so that they can participate in markets and system operation as a single entity [13]. VPPs are now operating commercially in several countries. MAEC's community hub resembles a VPP at the transformer or feeder scale; the distinctive question is how hub-level decisions respect local network limits while also serving higher layers.

3.4 Transactive energy

Transactive energy coordinates devices through economic and control signals exchanged among participants [14]. Reviews have described transactive control frameworks for systems with high shares of distributed resources [15], market-based coordination of distributed resources [16], integrated approaches combining home energy management, distributed optimal power flow and peer-to-peer trading [17], and mathematical models used in transactive systems [18]. Transactive approaches are a principal benchmark for MAEC, and Baseline B (Section 7) includes a transactive or market-based variant.

3.5 Network-aware and distributed optimisation

Distributed optimisation methods, including the alternating direction method of multipliers (ADMM) [21], have been applied widely to power systems [25]. Linearised distribution power flow models derived from the branch-flow formulation, often tested on the 33-bus system introduced by Baran and Wu [20], allow network constraints to be embedded in tractable optimisation. A comprehensive review of planning, modelling, optimisation and control of distributed energy systems is given in [19]. MAEC's solution approach (Section 5.10) builds on these methods rather than proposing a new optimisation algorithm.

3.6 Uncertainty and resilience

Robust optimisation [22], stochastic programming and model-predictive control are established tools for operating under forecast uncertainty. Resilience research addresses extreme weather, contingencies and restoration. MAEC's reliability formulation combines a consequence-weighted cost of unserved energy with tiered commitments; whether this improves outcomes relative to conventional stochastic or robust scheduling is an empirical question.

3.9 Preliminary novelty assessment

The review above, which is preliminary and must be completed systematically, suggests the following positioning:

Research element	Preliminary status	Evidence required
Five-layer hierarchy	Not novel in itself	None; acknowledged as prior art
Network-aware coordination of distributed resources	Substantial prior work	Comparison with strongest prior methods
Value allocation across more than two layers without double counting (M1)	Possibly under-addressed in combination with M2 and M3	Systematic review; formal properties; simulation
Bidirectional inter-layer guardrails (M2)	Related to constraint-passing in hierarchical and distributed optimisation	Precise specification distinguishing it from existing methods
Tiered, revisable commitments with defined revision penalties (M3)	Related to reserve products and probabilistic bidding	Comparison with stochastic bidding and reserve products
Joint M1 + M2 + M3 framework with consequence-weighted reliability	Candidate research opportunity	Ablation study showing added value

Originality test. Before novelty is claimed, it must be possible to state in two or three precise sentences how MAEC differs from the closest existing paper. Appendix B provides the novelty matrix template used for this purpose.

3.10 Search protocol and patent considerations

The literature review will search Google Scholar, IEEE Xplore, ScienceDirect and OpenAlex using, individually and in combination: "hierarchical energy

3.7 Value allocation and cooperative games

The Shapley value [23] allocates the value created by a coalition according to members' average marginal contributions and satisfies efficiency, symmetry and additivity. Cooperative game theory has been applied to energy communities and cost sharing; the computational cost of exact Shapley values grows exponentially with the number of players, motivating sampling and structural approximations. Ostrom's work on common-pool resources [24] provides institutional principles (clear boundaries, local rule-making, monitoring, graduated sanctions) relevant to the governance of shared transformer capacity.

3.8 India-specific studies

The IEA's study of renewables integration in India examined flexibility, storage, transmission, demand response and state differences [5]; its work on grid-interactive buildings addresses building flexibility [6]; and RMI has proposed a national grid flexibility programme [7]. These provide system context and potential case-study parameters.

management" distributed energy resources; "multi-level energy coordination" power systems; "cross-layer coordination" smart grid; "network-aware transactive energy" reliability; "multi-agent optimization" renewable energy grid; "distributed energy resources" value allocation; "energy community" distribution grid coordination; and "robust optimization" renewable energy coordination India. Priority will be given to 2021-2026 publications, with backward and forward citation

searches. A broad coordination concept is unlikely to be patentable; a specific technical implementation may be, depending on jurisdiction. Prior-art searches will use India's InPASS and WIPO PATENTSCOPE [30] across four groups (architecture, optimisation, grid interaction and implementation), reviewing claims, families, priority dates and legal status. If a potentially patentable mechanism is identified, professional advice will be sought before publishing enabling details. A patent search is not a legal opinion.

IV. ARCHITECTURE, INTERFACES AND COORDINATION RULES

4.1 Design principles

MAEC is guided by six principles:

1. **Subsidiarity.** Decisions are taken at the lowest layer that has the information and authority to take them; higher layers set limits and procure services rather than controlling devices directly.
2. **Owner primacy.** A device follows its owner's constraints (comfort bands, backup reserves, medical-priority flags, opt-out) before any external signal.
3. **Bidirectional guardrails.** Each layer publishes security limits to its neighbours, and each layer optimises only within the limits it has received.

4. **Graded commitments.** Inter-layer offers carry an explicit firmness level and can be revised under defined rules.
5. **Service-based remuneration.** Each distinct service is valued at the layer where it creates value and is paid once.
6. **Graceful degradation.** Loss of communication with a higher layer causes a layer to revert to safe, pre-agreed local behaviour rather than to stop operating.

4.2 The five layers



Figure 2. MAEC layers and the shared decision logic that connects them.

Layer	Scope	Decision rights	Principal objective	Typical timescale
L1 Device	Inverter, battery, charger, AC, pump, water heater	Local fast response within owner and safety limits	Follow schedule; protect device; respond to local frequency and voltage	Sub-second to minutes
L2 Building	Home, shop, farm, factory, public building	Schedule devices; set owner priorities	Owner cost and comfort or production subject to guardrails	Minutes to day-ahead
L3 Community hub	Buildings on one distribution transformer or feeder segment; village; cluster	Pool and share local resources; protect local network	Local network security and fair sharing; deliver commitments	15 minutes to day-ahead
L4 Distribution and state	DISCOM network and State Load Dispatch Centre	Procure flexibility; enforce feeder and substation limits	Network security and least-cost procurement	15 minutes to years
L5 Region and nation	Regional and national load dispatch; markets; planning	System balance, reserves and inter-state flows	System security and least-cost balancing	Seconds to years

4.3 Shared decision logic

Every layer implements six functions, scaled to its context: forecasting (load, renewable output, prices, equipment state); optimisation (its own schedule); uncertainty management (margins and scenarios); privacy (sharing aggregates rather than individual data); incentives (receiving and paying for services); and reliability guardrails (respecting received limits and publishing its own).

4.4 Interfaces and message types

Adjacent layers exchange a small set of standard messages. The set below is a candidate minimal vocabulary; it is intended to be mapped onto existing standards (Section 9.4) rather than replace them.

Message	Direction	Content	Example
Guardrail	Both	Upper and lower limits on net exchange or voltage per interval, with validity period	Hub to DISCOM: import at most 140 kW, export at most 110 kW
Flexibility request	Down	Quantity, interval, location, price or price curve	DISCOM to hubs on Feeder 17: reduce 400 kW, 20:00 to 21:00
Tiered offer	Up	Firm, likely and possible quantities per interval	Hub: 60 kW firm, 40 kW likely, 30 kW possible
Revision	Up	Changed tier quantities and reason code	Cooler forecast; two homes away
Dispatch	Down	Activated quantities within accepted offers	Activate firm plus likely for 20:00 to 21:00
Measurement	Up	Metered delivery relative to baseline	Delivered 92 kW reduction
Settlement	Down	Payments by service and layer	Statement lines for hub, DISCOM and system services
Emergency state	Down	Stress level and priority rules	Heatwave level 2: protect Tier 1 and Tier 2 loads

4.5 Mechanism M1: cross-layer value allocation

A single action, such as a battery discharging 3 kWh between 19:00 and 20:00, may simultaneously reduce the owner's peak-tariff purchases, avoid overloading a distribution transformer, relieve feeder congestion, and contribute to system ramping. M1 requires that (i) each service be valued as an avoided cost at the layer where it occurs, (ii) the sum of payments for a service not exceed the avoided cost to the paying layer, (iii) a benefit already reflected in a lower-layer price not be paid again at a higher layer, and (iv) where several resources jointly create a benefit, the payment be divided by a rule with defined fairness properties, such as the Shapley value or a tractable approximation (Section 5.9).

4.6 Mechanism M2: inter-layer reliability guardrails

Each layer publishes guardrails to its neighbours. A hub publishes its transformer's import and export limits and its committed backup reserve to the DISCOM; the DISCOM publishes feeder-level import caps or voltage bands to the hub; a building publishes its owner's minimum backup reserve and comfort bands to the hub. Each layer's optimisation treats received guardrails as hard constraints (Section 5.3). Guardrails are time-varying and may be tightened during declared stress states. The hypothesis is that explicit, bidirectional guardrails prevent the four failure modes in Section 2.2 at lower cost than either centralised control or price signals alone.

4.7 Mechanism M3: adaptive tiered commitments

A hub's offer to the layer above is divided into three tiers, analogous to confirmed, reservation-against-cancellation and waitlisted railway tickets: firm (delivered with high probability, highest payment), likely (delivered with moderate probability, lower payment) and possible (paid only if activated). Offers may be revised as forecasts improve, with revision penalties that rise as delivery approaches (Section 5.7). The hypothesis is that tiered, revisable commitments allow higher layers to count small resources more accurately than binary commitments, reducing both over-procurement of reserves and the frequency of unmet commitments.

4.8 Emergency operation and priority tiers

During declared stress states, the hub and DISCOM apply priority tiers (Figure 3). Tier 4 loads (shiftable charging, heating, pumping with storage, batch processes) are reduced first; Tier 3 loads (ordinary cooling within comfort limits, commercial loads) are reduced gently; Tier 2 loads (registered medical households, vaccine and food cold chains, essential lighting and fans) are strongly protected; and Tier 1 loads (hospitals, life support, critical water and emergency services) are interrupted only if physically unavoidable. During islanding, a hub's local critical loads take precedence over higher-layer requests.



Figure 3. Priority tiers used in the consequence-weighted reliability formulation.

4.9 Security and privacy by design

Devices must reject commands that violate built-in safety envelopes regardless of their source; must apply randomised response delays and rate limits to prevent synchronised switching; and must revert to safe local schedules if signals are lost or fail authentication. Large aggregated actions require confirmation through more than one channel. Personal data remain at L1 and L2 by default; hubs share only aggregates and commitments with L4; and data sharing beyond these

5.2 Planning objective

At the planning level, the objective is to minimise total annualised system cost subject to equivalent reliability and operating constraints:

$$\min J = C^{\text{gen}} + C^{\text{sto}} + C^{\text{net}} + C^{\text{op}} + C^{\text{use}} \quad (1)$$

where C^{gen} is annualised generation cost including fuel, C^{sto} annualised storage cost including degradation, C^{net} annualised network cost, C^{op} operating and coordination cost (including communication and software), and C^{use} the expected cost of unserved energy. Each cost must be counted exactly once; for example, battery degradation appears in C^{sto} and not again in C^{op} . The reliability term prevents the model from reducing cost simply by accepting more outages.

5.3 Layer subproblems with guardrails

Each layer l solves a subproblem over its own decision variables x_l :

$$\min_{x_l \in \mathcal{X}_l} f_l(x_l) \quad \text{s. t.} \quad g_l(x_l) \leq G_{l-1 \rightarrow l}, \quad h_l(x_l) \leq G_{l+1 \rightarrow l} \quad (2)$$

where f_l is the layer's local objective (owner cost and discomfort at L2; local security and fair sharing at L3; procurement and network cost at L4), and G denotes guardrails received from the layers below and above. For a hub, the aggregate exchange and the transformer guardrail are:

$$P_{h,t} = \sum_{b \in B_h} \sum_{d \in \mathcal{D}_b} p_{d,t}, \quad \underline{P}_{h,t}^G \leq P_{h,t} \leq \overline{P}_{h,t}^G \quad (3)$$

purposes is consent-based, consistent with India's Digital Personal Data Protection Act, 2023 [27].

V. MATHEMATICAL FORMULATION

5.1 Notation

Symbol	Meaning
$t, \Delta t$	Time interval and its length (0.25 h)
s, π_s	Scenario and its probability
d, b, h	Device, building, community hub
$p_{\{d,t\}}$	Net power of device d (positive = import)
$P_{\{h,t\}}$	Net import of hub h at its transformer
$e_{\{b,t\}}$	Battery state of energy
$\theta_{\{b,t\}}$	Indoor temperature of a thermally controlled zone
$U_{\{k,t,s\}}$	Unserved demand of priority tier k
VoLL_k	Value of lost load for tier k
G	Guardrail limits received from a neighbouring layer
F, L, O	Firm, likely and possible commitment quantities
ϕ_i	Value allocated to participant i

5.4 Network model

At L4, distribution network constraints are represented with the linearised branch-flow (LinDistFlow) model, in which v denotes squared voltage magnitude:

$$v_{j,t} = v_{i,t} - 2(r_{ij}P_{ij,t} + x_{ij}Q_{ij,t}), \quad \underline{v} \leq v_{j,t} \leq \bar{v}, \quad P_{ij,t}^2 + Q_{ij,t}^2 \leq \bar{S}_{ij}^2 \quad (4)$$

Transformer thermal limits are represented as apparent-power limits, optionally with time-dependent ratings reflecting ambient temperature.

5.5 Storage and degradation

Battery state of energy evolves with charging and discharging efficiencies and is bounded below by the owner's committed backup reserve:

$$e_{b,t+1} = e_{b,t} + (\eta^c p_{b,t}^c - p_{b,t}^d / \eta^d) \Delta t, \quad e_{b,t}^{\text{res}} \leq e_{b,t} \leq \bar{E}_b \quad (5)$$

A first-order throughput-based degradation cost per unit of delivered energy is used, consistent with approaches that factor cycle ageing into market participation [31]:

$$c_b^{\text{deg}} = \frac{C_b^{\text{inv}}}{N_b^{\text{cyc}} \bar{E}_b \eta_b^{\text{rt}}}, \quad C_b^{\text{deg}} = c_b^{\text{deg}} \sum_t p_{b,t}^d \Delta t \quad (6)$$

For example, a 5-kWh battery costing ₹2,50,000 with 4,000 cycles and 90% round-trip efficiency has c^{deg} of about ₹13.9 per kWh delivered, which exceeds typical residential tariff differentials. This illustrates why M1's multi-service remuneration may be decisive for the economics of distributed storage.

5.6 Thermal and shiftable loads

Air-conditioned zones are represented by a first-order thermal model with comfort bands set by the owner:

$$\theta_{t+1} = a \theta_t + (1-a)(\theta_t^{\text{out}} - R \eta^{\text{ac}} p_t^{\text{ac}}), \quad \underline{\theta}_b \leq \theta_t \leq \bar{\theta}_b \quad (7)$$

Shiftable loads (pumping with storage, water heating, charging) are represented by energy requirements within time windows; cold rooms and ice storage use analogous thermal models.

5.7 Tiered commitments and revision penalties

A hub's offer in interval t comprises firm F , likely L and possible O quantities, with chance constraints of decreasing stringency:

$$\Pr(y_{h,t} \geq F_{h,t}) \geq 1 - \epsilon^F, \quad \Pr(y_{h,t} \geq F_{h,t} + L_{h,t}) \geq 1 - \epsilon^L, \quad \epsilon^F < \epsilon^L \quad (8)$$

Payment combines tier prices, activation of possible offers, and a shortfall penalty on firm commitments:

$$\Pi_{h,t} = \rho^F F_{h,t} + \rho^L L_{h,t} + \rho^O y_{h,t}^O - \kappa^F [F_{h,t} - y_{h,t}]^+ \quad (9)$$

Revisions made with lead time τ before delivery incur a penalty that increases as τ decreases:

$$C_{h,t}^{\text{rev}} = \kappa^{\text{rev}}(\tau) |\Delta F_{h,t}|, \quad \frac{d\kappa^{\text{rev}}}{d\tau} < 0 \quad (10)$$

The design question for research is how to set ρ , κ and ϵ so that honest reporting is incentivised and participation is not discouraged.

5.8 Consequence-weighted unserved energy

Rather than a single value of lost load, unserved energy is weighted by priority tier k :

$$C^{\text{use}} = \sum_s \pi_s \sum_t \sum_{k=1}^4 \text{VoLL}_k U_{k,t,s} \Delta t, \quad \text{VoLL}_1 \gg \text{VoLL}_2 > \text{VoLL}_3 > \text{VoLL}_4 \quad (11)$$

with $U \in \{1, t, s\}$ constrained to zero except where physically infeasible. Tier values should be informed by survey or revealed-preference studies and treated as policy parameters subject to sensitivity analysis.

5.9 Value allocation (M1)

For a coalition N of participants jointly creating value $v(N)$, the Shapley allocation is:

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|! (n - |S| - 1)!}{n!} [v(S \cup \{i\}) - v(S)] \quad (12)$$

It satisfies efficiency (allocations sum to $v(N)$), symmetry and additivity. Two further conditions are imposed to prevent double counting across layers: the value of each service σ is defined as the avoided cost at the layer where it accrues, measured against a counterfactual in which other services are unchanged; and the total paid by each layer cannot exceed its avoided cost:

$$\sum_i \phi_i = v(N), \quad \sum_i \text{Pay}_{i,\sigma} \leq \Delta C_\sigma \quad \forall \sigma \in \{\text{own}, \text{hub}, \text{DSO}, \text{sys}\} \quad (13)$$

Because exact computation is exponential in n , MAEC will test a permutation-sampling estimator, with M sampled orderings and predecessor set P_i^m :

$$\hat{\phi}_i = \frac{1}{M} \sum_{m=1}^M [v(P_i^m \cup \{i\}) - v(P_i^m)] \quad (14)$$

together with a hierarchical approximation in which Shapley values are computed among hubs at L4 and then among buildings within each hub at L3. The research questions are the approximation error, computational time and participants' perceived fairness.

5.10 Distributed solution approach

The coupled layer problems are coordinated using ADMM-type consensus [21], in which hubs and the DISCOM iterate on the hub exchange P_h (Equations 15 to 17):

$$x_h^{k+1} = \arg \min_{x_h} f_h(x_h) + \frac{\rho}{2} \|P_h(x_h) - z_h^k + u_h^k\|_2^2 \quad (15)$$

$$z^{k+1} = \arg \min_{z \in Z_G} f_{\text{DSO}}(z) + \frac{\rho}{2} \sum_h \|P_h(x_h^{k+1}) - z_h + u_h^k\|_2^2 \quad (16)$$

$$u_h^{k+1} = u_h^k + P_h(x_h^{k+1}) - z_h^{k+1} \quad (17)$$

where Z_G is the DISCOM's feasible set including network constraints and guardrails received from L5. Only aggregate hub exchanges are communicated, which bounds message size per interval by the number of hubs rather than the number of devices. Convergence speed, communication volume and behaviour under message loss are evaluation metrics (Section 7.4). Device-level fast response at L1 is implemented with local droop rules and does not depend on this iteration.

VII. INDIAN CASE-STUDY SYSTEM, DATASETS AND ASSUMPTIONS

6.1 Test systems

Three test systems are proposed, chosen to represent distinct Indian contexts while remaining reproducible:

Test system	Description	Primary questions
T1 Urban low-voltage network	One 100 to 250 kVA distribution transformer serving 60 to 150 households with rooftop solar, air conditioners, water pumps, electric two-wheelers and some batteries; optionally an apartment complex	Transformer overload, rooftop overvoltage, cooling peaks, herding; M2 and M1 at L2 and L3
T2 Rural agricultural feeder	An 11 kV feeder with feeder-level solar, 200 to 400 irrigation pumps, a cold room, a milk chiller and a village load	Daytime pumping, productive loads as flexibility, farmer remuneration, groundwater-aware scheduling
T3 Mixed medium-voltage network	An adapted IEEE 33-bus feeder [20] with several T1-type hubs and one T2-type hub, connected to a state-level system representation	L3 to L4 coordination, M3 commitments, ADMM performance, stress scenarios

6.2 Data sources

Data	Potential sources	Notes
Feeder and transformer parameters, loading, failure records	Participating DISCOMs	Access conditions to be confirmed; anonymisation required
Smart-meter interval data	DISCOMs; pilot participants with consent	Aggregated where possible
State demand and generation	State Load Dispatch Centres; Grid-India	Public reports and requested data
Planning statistics	Central Electricity Authority	Public reports
Solar irradiance and weather	Public meteorological and satellite datasets	Multi-year series for stress scenarios
Market prices	Indian power exchanges	Public day-ahead and real-time price series
Appliance and pump characteristics	Manufacturer data; field surveys	To calibrate flexibility models

Data availability, granularity and access conditions will be confirmed before the test systems are finalised. Where real data are unavailable, synthetic profiles calibrated to published statistics will be used and labelled as such. Results will always distinguish simulated, laboratory and field evidence.

6.3 Key assumptions

Assumption	Base value or approach	Sensitivity range
Participation rate of households	30%	10% to 70%
Air-conditioner flexibility	Comfort band of plus or minus 1 °C around owner setpoint	0.5 °C to 2 °C
Battery cost and life	Section 5.5 example values	Cost minus 50% to plus 20%; 2,000 to 6,000 cycles
Forecast error (day-ahead solar)	Empirical distribution from historical data	Scaled 0.5 to 2 times
Communication availability	99% of intervals	90% to 100%
Values of lost load by tier	Policy parameters	Factor of 0.25 to 4
Time-of-day tariff spread	Per applicable state tariff order	0 to 2 times

VII. BASELINES, EXPERIMENTS AND EVALUATION METRICS

7.1 Compared methods

Method	Description	Purpose
Baseline A	Existing resources operated conventionally: fixed tariffs, uncoordinated rooftop export, no flexibility procurement	Reference for current practice
Baseline B1	Centralised network-aware stochastic optimisation with full information	Performance bound against which MAEC's efficiency loss is measured
Baseline B2	Established decentralised method: price-based transactive coordination or a VPP optimised for system services	Strong practical competitor
MAEC	Same resources and constraints coordinated with M1, M2 and M3	Proposed method

All methods use identical resources, input data, network models and reliability requirements. Any improvement must therefore come from coordination, not from additional investment or more conservative operating assumptions.

7.2 Ablation design

To attribute performance to individual mechanisms, MAEC will be run in seven configurations: each mechanism alone (M1, M2, M3), each pair (M1+M2, M1+M3, M2+M3) and all three together. Comparing these configurations with Baseline B2 identifies which mechanisms add value and whether they interact.

7.3 Scenario set

Scenario	Description
S1 Normal operation	Representative days across seasons
S2 Heatwave	Five consecutive days with elevated temperatures and evening and night demand
S3 Low-renewable period	Multi-day monsoon cloud cover or winter fog with reduced solar output, and low-wind days
S4 Contingency	Loss of a large generator (frequency event) and loss of a feeder or transformer
S5 Communication failure	Loss or corruption of inter-layer messages for defined periods
S6 Herding stress	Synchronised tariff boundary or common signal to many devices
S7 Cyclone and islanding	Upstream outage of several days with hub islanding and restoration

Stress scenarios will be constructed from historical weather records where possible, so that multi-day events are realistic rather than assumed.

7.4 Metrics

Metric	Definition	Establishes
Total annualised system cost	J in Equation (1)	Economic competitiveness
Expected energy not served	Sum over scenarios of unserved energy, reported by priority tier	Reliability and protection of critical loads
Renewable curtailment	Curtailed energy divided by available renewable energy	Use of clean generation
Peak loading and congestion	Maximum and 95th-percentile transformer and feeder loading; hours above rating	Asset utilisation and life
Voltage compliance	Share of node-intervals within statutory voltage limits	Power quality
Emissions	Grid emissions using marginal or average factors, reported separately	Environmental performance
Computation and communication	Solve time per interval; messages and bytes per interval; convergence iterations	Scalability
Recovery	Time to restore served load after disruption; load served during islanding	Resilience
Fairness of outcomes	Distribution of net benefit across households and income groups; Jain's index	Equity and incentives

Jain's fairness index for net benefits x_i across n participants [32] is:

$$\mathcal{J}(x) = \frac{\left(\sum_{i=1}^n x_i\right)^2}{n \sum_{i=1}^n x_i^2}, \quad \frac{1}{n} \leq \mathcal{J} \leq 1 \quad (18)$$

It will be complemented by reporting benefits for the lowest-income and non-owning participants

explicitly, since an index alone can conceal harm to small groups.

7.5 Statistical treatment

Results will be reported as distributions over scenarios and random seeds, with confidence intervals. Differences between methods will be tested with paired comparisons on identical scenarios. All code, parameters and synthetic data will be published to allow reproduction.




Figure 4. Evaluation design: identical inputs, four methods, seven ablation configurations and seven scenario families, scored on nine metrics.

VIII. HYPOTHESES, ILLUSTRATIVE EXAMPLES AND EXPECTED RESULTS

8.1 Falsifiable hypotheses

No empirical results are reported in this version. The following hypotheses will be tested:

ID	Hypothesis	Rejected if
H1	MAEC reduces total system cost relative to Baseline B2 at equal reliability	Cost difference is not significant or favours B2
H2	M2 guardrails reduce transformer overload hours and voltage violations relative to B2 without increasing cost by more than the avoided asset damage	No reduction, or cost increase exceeds avoided damage
H3	M3 tiered commitments reduce both unmet commitments and procured reserve relative to binary commitments	Neither quantity improves
H4	M1 allocation increases participation-weighted flexibility at equal total payment relative to single-service remuneration	No increase in delivered flexibility
H5	Consequence weighting reduces Tier 1 and Tier 2 unserved energy in stress scenarios at small or no increase in total unserved energy	Priority loads are not better protected
H6	MAEC's efficiency loss relative to centralised B1 is small (for example, within a few per cent of cost) while communication volume scales with hubs rather than devices	Efficiency loss is large or communication does not scale

A finding that MAEC does not outperform established methods would itself be a useful result, indicating that existing approaches are adequate or that only a subset of mechanisms adds value.

8.2 Illustrative example 1: cross-layer value allocation

Consider three participants on one transformer during a single evening peak: A, a group of smart air conditioners in a cooling rotation; B, a household battery; and C, a community battery. Their joint action avoids transformer overload, reduces feeder congestion and lowers peak purchase cost. Suppose the avoided costs of each coalition, in rupees for that evening, are as follows. These numbers are illustrative only.

Coalition S	$v(S)$ (₹)
{A}	200
{B}	300
{C}	400
{A, B}	700
{A, C}	800
{B, C}	900
{A, B, C}	1,500

The grand coalition is worth more than the sum of individuals because transformer damage is avoided only when the combined reduction is large enough. Applying Equation (12) over all six orderings gives Shapley allocations of ₹400 for A, ₹500 for B and ₹600 for C, which sum to ₹1,500 (efficiency). By comparison, paying each participant only its stand-

alone value would distribute ₹900 and leave ₹600 of jointly created value unrecognised, weakening the incentive to participate. The full calculation is given in Appendix C.

8.3 Illustrative example 2: transformer guardrail

A 100 kVA transformer at an assumed power factor of 0.9 has a real-power capacity of about 90 kW. With 25 air conditioners at 1.5 kW and 70 households at an average 0.8 kW of other load, evening demand is about 93.5 kW, or about 104% of capacity. If a guardrail of 85 kW is published for the peak interval and the hub rotates one-third of air conditioners into a reduced state (about 12.5 kW), loading falls to about 81 kW, or 90%. Under Baseline B2, a system-level price signal that does not observe this transformer would provide no assurance that loading falls below the limit, and could increase it if devices respond to a common tariff boundary.

8.4 Expected outputs

Expected outputs of the first research phase are: a formal specification; open-source simulation code for T1 to T3; results tables for all methods, ablations and scenarios; a sensitivity analysis; and a short policy brief summarising findings and limitations.

IX. ECONOMIC, REGULATORY AND OPERATIONAL FEASIBILITY

9.1 Economic feasibility

MAEC competes with conventional investments: transformer and feeder upgrades, peaking capacity and emergency power purchases. An illustrative comparison for a DISCOM considering upgrading 100 overloaded transformers at ₹6 lakh each (₹6 crore) shows that community hubs costing ₹1.5 lakh each to establish plus ₹50,000 per transformer per year in flexibility payments would cost ₹4 crore over five years, but ₹6.5 crore if payments were ₹1 lakh per year. Flexibility is therefore not automatically cheaper; it must be priced and compared case by case, which the planning formulation in Equation (1) is designed to do.

9.2 Tariff design

Time-of-day tariffs, provided for in the Electricity (Rights of Consumers) Rules, 2020 as amended [26], are a prerequisite for price-responsive flexibility. To

protect small consumers, the design should include an essential-use block priced uniformly across the day, revenue neutrality for typical small users, and access to participation credits for households without rooftop systems or batteries, including through community solar.

9.3 Regulatory and institutional requirements

Requirement	Possible instrument
Legal recognition of community hubs	State regulatory provisions defining hub operators (DISCOM, cooperative, resident welfare association, panchayat or licensed aggregator), rights and obligations
Flexibility procurement by DISCOMs	Feeder-level flexibility programmes with transparent pricing and measurement
Multi-service remuneration without double counting	Settlement rules implementing M1 conditions
Consumer protection	Opt-out rights, comfort guarantees, essential-use protection, dispute resolution
Testing before rule-making	Regulatory sandboxes with independent evaluation
Data protection	Consent-based data sharing consistent with [27]

9.4 Standards and interoperability

MAEC's message vocabulary (Section 4.4) should be mapped onto existing standards for distributed-resource communication and demand response, such as IEEE 2030.5, OpenADR and IEC 61850 [29], rather than creating a parallel standard. A voluntary "flexibility-ready" appliance specification for air conditioners, chargers, pumps and water heaters would avoid locking in non-responsive devices, since appliances sold today remain in service for a decade or more.

9.5 Operational and lifecycle considerations

Operational feasibility depends on local maintenance capacity, especially in rural areas, where field reports indicate that equipment such as solar pumps can remain idle when local technicians are unavailable [4]. Explicit degradation costs (Section 5.5) and guardrails that prevent chronic transformer overload are expected to extend asset life. End-of-life obligations under India's Battery Waste Management Rules, 2022 and E-Waste (Management) Rules, 2022 [28] should be included in hub business models, together with standards for testing and reusing second-life batteries.

9.6 Productive-use extension

In rural feeders (test system T2), the hub objective can be extended to include productive outcomes: reduced post-harvest losses through cold storage, milk chilling, food processing and daytime irrigation. Evaluation must distinguish genuine value creation (reduced waste, additional output, reduced fuel use) from transfers between participants (for example, between farmers and traders), and must account for groundwater impacts of pumping.

X. LIMITATIONS, THREATS TO VALIDITY AND ROADMAP

10.1 Limitations

This document presents a design and hypotheses, not results. The novelty assessment is preliminary. The mathematical formulation uses linearised network models and first-order thermal and degradation models, which may not capture all relevant physics. Values of lost load by priority tier are policy parameters with substantial uncertainty. Behavioural responses, including participation, opt-out and reporting honesty under M3, are assumed rather than measured.

10.2 Threats to validity

Threat	Mitigation
Improvement due to extra resources rather than coordination	Identical resources and constraints across methods (Section 7.1)
Improvement due to conservatism rather than intelligence	Report reserve held and cost of conservatism; compare at equal reliability
Weak baseline	Include centralised bound (B1) and strong decentralised method (B2)
Optimistic flexibility assumptions	Calibrate from field data; wide sensitivity ranges
Perfect communication assumed	Communication-failure scenario S5
Synthetic data unrepresentative	Validate synthetic profiles against published statistics; label clearly
Results specific to one network	Three test systems; sensitivity to topology and penetration

10.3 Twelve-month roadmap

Months	Activities	Deliverables
1 to 2	Systematic literature review; novelty matrix; patent search; precise problem definition	Concept note; preliminary novelty assessment

3 to 4	Formal specification of layers, messages, rules, objectives and constraints	Architecture specification; simulation design
5 to 7	Data acquisition; implementation of baselines and MAEC; experiments; sensitivity analysis	Reproducible model; first results
8 to 9	Analysis and manuscript preparation; expert feedback	Journal manuscript; updated white paper with results
10 to 12	Journal submission; pilot partner engagement; policy brief	Submission package; pilot proposal

This schedule is indicative; data access, collaborators and funding may change it substantially.

10.4 Publication and engagement strategy

A two-track approach is proposed. The technical track targets peer-reviewed journals in renewable energy and power systems research, such as Renewable Energy and Electric Power Systems Research, after results are available; acceptance is not assumed. The policy track prepares a shorter brief for DISCOMs, regulators, research institutions and potential pilot partners, focused on the problem, the proposal, evidence to date and a validation plan. A first field pilot at a single distribution transformer, run over a full summer with a cooperating DISCOM and an academic partner, is recommended before any scaled programme.

XI. CONCLUSION

India's renewable expansion has made coordination, rather than capacity alone, a central constraint. MAEC proposes a five-layer architecture and three coordination mechanisms intended to let distributed resources serve their owners, their local networks and the wider system without conflict and with fair remuneration. The proposal deliberately builds on established work and is framed as a set of falsifiable hypotheses. Its value will be determined by careful comparison against strong baselines under realistic and stressful conditions. If the hypotheses hold, the mechanisms could inform tariff design, community-hub regulation and appliance standards in India and in other renewable-dominant systems; if they do not, the evaluation will clarify which existing approaches are sufficient.

REFERENCES

- [1]. Ministry of New and Renewable Energy, Government of India, "Installed renewable energy capacity as of June 30, 2026," as reported in "India's Installed Renewable Energy Capacity Reaches 288.58 GW as Solar Soars," Angel One News, 2026.
- [2]. Down To Earth, "India Adds Record 29 GW Solar, Wind in First Half of 2026," 2026.
- [3]. Mercom India, "PM Surya Ghar Powers 5 million Homes, Total Capacity Hits 15 GW," 2026.
- [4]. "PM-KUSUM 2 Set to Launch Soon, Announces Union Minister Pralhad Joshi," NewKerala, Aug. 2026.
- [5]. International Energy Agency, Renewables Integration in India. Paris, France: IEA.
- [6]. International Energy Agency, Efficient Grid-Interactive Buildings in India. Paris, France: IEA.
- [7]. RMI, Envisioning a National Grid Flexibility Programme for India.
- [8]. P. Denholm, M. O'Connell, G. Brinkman, and J. Jorgenson, Overgeneration from Solar Energy in California: A Field Guide to the Duck Chart. National Renewable Energy Laboratory, Tech. Rep. NREL/TP-6A20-65023, 2015.
- [9]. R. H. Lasseter, "MicroGrids," in Proc. IEEE Power Engineering Society Winter Meeting, 2002.
- [10]. D. E. Olivares et al., "Trends in microgrid control," IEEE Transactions on Smart Grid, vol. 5, no. 4, pp. 1905–1919, 2014.
- [11]. D. S. Callaway and I. A. Hiskens, "Achieving controllability of electric loads," Proceedings of the IEEE, vol. 99, no. 1, pp. 184–199, 2011.
- [12]. M. H. Albadi and E. F. El-Saadany, "A summary of demand response in electricity markets," Electric Power Systems Research, vol. 78, no. 11, pp. 1989–1996, 2008.
- [13]. D. Pudjianto, C. Ramsay, and G. Strbac, "Virtual power plant and system integration of distributed energy resources," IET Renewable Power Generation, vol. 1, no. 1, pp. 10–16, 2007.
- [14]. K. Kok and S. Widergren, "A society of devices: Integrating intelligent distributed resources with transactive energy," IEEE Power and Energy Magazine, vol. 14, no. 3, pp. 34–45, 2016.
- [15]. J. Hu, G. Yang, K. Kok, Y. Xue, and H. W. Bindner, "Transactive control: A framework for operating power systems characterized by high penetration of distributed energy resources," Journal of Modern Power Systems and Clean Energy, vol. 5, no. 3, pp. 451–464, 2017, doi: 10.1007/s40565-016-0228-1.
- [16]. "Towards a transactive energy system for integration of distributed energy resources: Home energy management, distributed optimal power flow, and peer-to-peer energy trading," Renewable and Sustainable Energy Reviews, vol. 132, 2020.
- [17]. "Review of transactive energy mathematical models," Energy Conversion and Economics, 2023, doi: 10.1049/enc2.12076.
- [18]. "A comprehensive review of planning, modeling, optimization, and control of distributed energy systems," Carbon Neutrality, 2022, doi: 10.1007/s43979-022-00029-1.
- [19]. M. E. Baran and F. F. Wu, "Network reconfiguration in distribution systems for loss reduction and load balancing," IEEE Transactions on Power Delivery, vol. 4, no. 2, pp. 1401–1407, 1989.
- [20]. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," Foundations and Trends in Machine Learning, vol. 3, no. 1, pp. 1–122, 2011.
- [21]. D. Bertsimas and M. Sim, "The price of robustness," Operations Research, vol. 52, no. 1, pp. 35–53, 2004.
- [22]. L. S. Shapley, "A value for n-person games," in Contributions to the Theory of Games II, H. W. Kuhn and A. W. Tucker, Eds. Princeton, NJ, USA: Princeton University Press, 1953, pp. 307–317.
- [23]. E. Ostrom, Governing the Commons: The Evolution of Institutions for Collective Action. Cambridge, U.K.: Cambridge University Press, 1990.
- [24]. D. K. Molzahn et al., "A survey of distributed optimization and control algorithms for electric power systems," IEEE Transactions on Smart Grid, vol. 8, no. 6, pp. 2941–2962, 2017.

- [25]. Ministry of Power, Government of India, Electricity (Rights of Consumers) Rules, 2020, as amended.
- [26]. Government of India, Digital Personal Data Protection Act, 2023.
- [27]. Ministry of Environment, Forest and Climate Change, Government of India, Battery Waste Management Rules, 2022; E-Waste (Management) Rules, 2022.
- [28]. IEEE, IEEE Standard 2030.5: Smart Energy Profile Application Protocol; OpenADR Alliance, OpenADR 2.0 Specifications; International Electrotechnical Commission, IEC 61850: Communication Networks and Systems for Power Utility Automation.
- [29]. Office of the Controller General of Patents, Designs and Trade Marks, InPASS Public Search. Available: <https://iprsearch.ipindia.gov.in/PublicSearch>; World Intellectual Property Organization, PATENTSCOPE.
- [30]. B. Xu, J. Zhao, T. Zheng, E. Litvinov, and D. S. Kirschen, "Factoring the cycle aging cost of batteries participating in electricity markets," IEEE Transactions on Power Systems, vol. 33, no. 2, pp. 2248–2259, 2018.
- [31]. R. Jain, D.-M. Chiu, and W. Hawe, A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer Systems. DEC Research Rep. TR-301, 1984.

APPENDIX A. GLOSSARY OF TERMS

Term	Meaning in this paper
Community hub	L3 coordinator for buildings sharing a distribution transformer, feeder segment, village or cluster
DISCOM	Distribution company responsible for the distribution network and retail supply
Guardrail	Time-varying limit published by one layer to a neighbouring layer and treated as a hard constraint
Tiered commitment	Offer divided into firm, likely and possible quantities with different reliability and payment
Value of lost load	Cost to consumers of unserved energy, here differentiated by priority tier

Ablation	Running the method with mechanisms selectively removed to attribute performance
Herdng	Synchronised response of many devices to a common signal, producing a new peak

APPENDIX B. NOVELTY MATRIX TEMPLATE

One row is completed for each closely related paper or project.

Field	What to record
Reference	Authors, year, DOI or URL
Architecture	Centralised, hierarchical, distributed or market-based
Coordination	How decisions and commitments are made
Network constraints	Voltage, congestion, thermal limits, security
Uncertainty	Forecast errors, outages, weather and demand
Economic treatment	Costs, incentives, allocation of benefits across layers
Reliability	Metrics and constraints used
Validation	Simulation, laboratory, field trial, test network or real data
Limitation	Explicitly reported or demonstrated limitation
Difference from MAEC	The exact feature (M1, M2, M3 or their combination) not already addressed

APPENDIX C. SHAPLEY CALCULATION FOR ILLUSTRATIVE EXAMPLE 1

For three participants there are six orderings. Each entry is the marginal contribution of the participant when it joins the coalition formed by those before it, using the values in Section 8.2.

Ordering	A	B	C
A, B, C	200	500	800
A, C, B	200	700	600
B, A, C	400	300	800
B, C, A	600	300	600
C, A, B	400	700	400
C, B, A	600	500	400
Sum	2,400	3,000	3,600
Shapley value (sum ÷ 6)	400	500	600

The allocations sum to ₹1,500, equal to $v(\{A, B, C\})$, confirming efficiency.