

Article

Not peer-reviewed version

---

# Coordinating Vehicle-to-Grid and Distributed Energy Resources in Multi-Dwelling Developments: A Real-Time Gateway Control Framework

---

[Janak Nambiar](#), [Samson Yu](#)<sup>\*</sup>, Ian Lilley, [Jag Makam](#), [Hieu Trinh](#)

Posted Date: 26 March 2026

doi: 10.20944/preprints202603.2085.v1

Keywords: vehicle-to-grid; virtual power plant; distributed energy resources; demand-side ancillary services



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Coordinating Vehicle-to-Grid and Distributed Energy Resources in Multi-Dwelling Developments: A Real-Time Gateway Control Framework

Janak Nambiar<sup>1</sup>, Samson Yu<sup>1,\*</sup>, Ian Lilley<sup>2</sup>, Jag Makam<sup>2</sup> and Hieu Trinh<sup>1</sup>

<sup>1</sup> School of Engineering, Deakin University, Waurn Ponds, Geelong, VIC 3216, Australia.

<sup>2</sup> Zeco Energy, 6c/148 Chesterville Rd, Cheltenham, VIC 3192, Australia.

\* Correspondence: samson.yu@deakin.edu.au

## Abstract

This study proposes a three-layer gateway control framework for a behind-the-meter virtual power plant (VPP) comprising vehicle-to-grid (V2G) capable electric vehicle (EV) chargers, battery energy storage systems (BESS), and rooftop photovoltaic (PV) generation in multi-dwelling residential developments. In particular, the first layer performs day-ahead scheduling to determine the hourly grid import baseline and frequency regulation ancillary service capacity for the following day. In the second layer, real-time regulation dispatch is performed by following the dynamic regulation signal from the grid operator, wherein V2G-capable EVs are coordinated alongside BESS as active demand-side participants in frequency regulation ancillary services, enabling the aggregated behind-the-meter fleet to respond to regulation signals in real time. The third layer performs per-minute three-phase load balancing to maintain network power quality compliance across the multi-dwelling site. The overall goal is to coordinate distributed energy resources behind a single network connection point to simultaneously reduce peak demand, maximise renewable self-consumption, and provide demand-side frequency regulation as a dispatchable VPP asset.

**Keywords:** vehicle-to-grid; virtual power plant; distributed energy resources; demand-side ancillary services

## 1. Introduction

Australia holds one of the world's highest per-capita rates of residential rooftop PV adoption, with installed capacity in the Victorian NEM exceeding 3.8 GW at the close of 2024 [1]. Concurrently, the Clean Energy Regulator projects residential PV capacity will exceed 6 GW by 2028 as module prices continue to fall below \$0.25/Wp. Electric vehicle registrations in Victoria grew 62% year-on-year during 2023-2024 [2], driven by the Zero Emissions Vehicle Subsidy and a widening range of models capable of bidirectional (V2G) charging. Multi-dwelling developments — apartment buildings, townhouse clusters, and mixed-use precincts — increasingly host all three asset classes on a shared network metering installation (NMI): rooftop PV, communal BESS, and V2G-capable EV charging bays.

Managing these co-located assets presents a hierarchical optimisation challenge spanning three timescales. At the day-ahead (24-hour) timescale, the site operator must determine how much energy to import, export, store, and bid into the AEMO demand-side frequency control ancillary service (D-FCAS) market while honouring time-of-use tariff windows, per-EV departure SoC requirements, and BESS cycle-life economics. At the five-minute real-time timescale, the controller must track the AEMO normalised RegD signal  $\omega(t) \in [-1, +1]$  with sufficient accuracy to maintain FCAS enablement despite stochastic EV availability, cloud-induced PV transients, and variable building load. At the one-minute phase-balancing timescale, the three-phase connection must remain within the AS/NZS

61000.3.3 current-imbalance limit, which is routinely violated when multiple single-phase EV chargers connect simultaneously on the same phase.

Existing literature reflects this fragmentation. Residential energy management systems are typically designed for single-family homes without ancillary service participation [3,4]. V2G coordination papers focus on grid-scale aggregation without addressing phase balance or NEM-specific demand-side frequency control ancillary service (D-FCAS) requirements [5,6]. FCAS participation studies concern utility-scale BESS and do not address the behind the meter (BTM) multi-dwelling context [7,8]. Phase-balancing works consider EV charger assignment in isolation without joint scheduling [9,10]. No prior work has integrated all four capabilities — day-ahead scheduling, V2G price-responsive dispatch, real-time D-FCAS regulation cascade, and AS/NZS three-phase balancing — in a single deployable gateway validated against Australian NEM conditions on embedded hardware.

This paper closes that gap with five specific contributions:

1. A corrected MILP formulation with import/export decomposition and terminal departure SoC constraints that eliminate all infeasible scheduling cycles.
2. An ARMA( $p, q$ ) forecasting module with Akaike Information Criterion (AIC) model-order selection and rolling bias correction, adopted and modified from [11].
3. A three-stage real-time cascade achieving 92.1-94.7% D-FCAS accuracy.
4. A greedy EV phase-reassignment algorithm reducing mean phase imbalance from 11.8% to 5.6%.
5. An open-source MATLAB implementation validated on the Raspberry Pi 4 with fully documented hardware interfaces. The controller is evaluated over seven days across three Melbourne field trial sites using Monte Carlo synthetic profiles calibrated against Victorian metered data.

The paper is organised as follows. Section 2 reviews related work across four sub-areas. Section 3 presents the site model and hardware architecture. Section 4 develops the ARMA forecasting methodology. Section 5 formulates the day-ahead scheduling in full detail. Section 6 describes the real-time regulation cascade. Section 7 presents the V2G dispatch logic. Section 8 describes the phase-balancing algorithm. Section 9 presents simulation results, sensitivity analysis, and benchmarking against prior work. Section 10 concludes the paper.

## 2. Related Work

### 2.1. Residential Energy Management Systems

Residential energy management systems optimising BESS and PV dispatch have attracted substantial attention. [3] proposed a real-time Markov decision process for single-home BESS dispatch under dynamic pricing, achieving 22% cost reduction.[4] applied a game-theoretic Nash equilibrium framework to a multi-home microgrid sharing a PV array, demonstrating convergence to a socially optimal equilibrium. Wang et al. [8] formulated a MILP for 12-apartment load scheduling in Sydney, reporting 31% peak demand reduction; however, neither EV V2G flexibility nor ancillary service participation were modelled. A comprehensive review by Mayer and Grof [12] of 84 peer-reviewed residential energy management system (REMS) papers found that fewer than 12% addressed multi-dwelling shared infrastructure and fewer than 5% validated on actual hardware — gaps the present work addresses directly.

Model predictive control (MPC) has emerged as a popular REMS framework because it naturally handles rolling forecast updates. However, incorporating binary V2G mode decisions and BESS mutual-exclusion constraints requires a mixed-integer extension, increasing the computational burden substantially. The mixed integer linear programming-plus-cascade architecture proposed here separates the binary scheduling problem — solved once per night within a generous 120-second time budget — from the real-time dispatch, executed as a linear cascade in milliseconds, making the

approach ran on a Raspberry Pi 4 without sacrificing the ability to model combinatorial asset interactions.

## 2.2. V2G Coordination

The study in [13] established the economic foundations of V2G, while [5] formulated V2G aggregation as a quadratic programme minimising distribution losses while maximising ancillary revenue. For multi-dwelling residential tablesettings, [14] studied OCPP-based EV coordination across a shared parking facility, demonstrating significant load-shifting potential, but did not address FCAS bid optimisation, BESS co-dispatch, or phase balance. A persistent modelling error in the V2G literature is the enforcement of EV departure SoC as a binding per-slot lower bound on stored energy, which prevents the scheduler from using EV energy early in the day even when spot prices are high. This paper instead enforces departure SoC as a terminal-only constraint at the departure slot, recovering 23% additional V2G dispatch capacity.

## 2.3. FCAS and Ancillary Service Participation

The authors [7] demonstrated grid-scale BESS frequency regulation in the NEM at 94% accuracy. [8] co-optimised energy arbitrage and FCAS bids for utility-scale storage, quantifying 23-38% revenue uplift from joint optimisation. The most directly related residential-scale work, [15], proposed an MPC scheme for BTM BESS D-FCAS participation reporting 91.6% regulation accuracy – close to the AEMO 92% threshold – without V2G or multi-phase constraints. The present work achieves 92.1-94.7% while adding V2G dispatch and phase balancing, using a computationally lighter cascade that is viable on commodity embedded hardware.

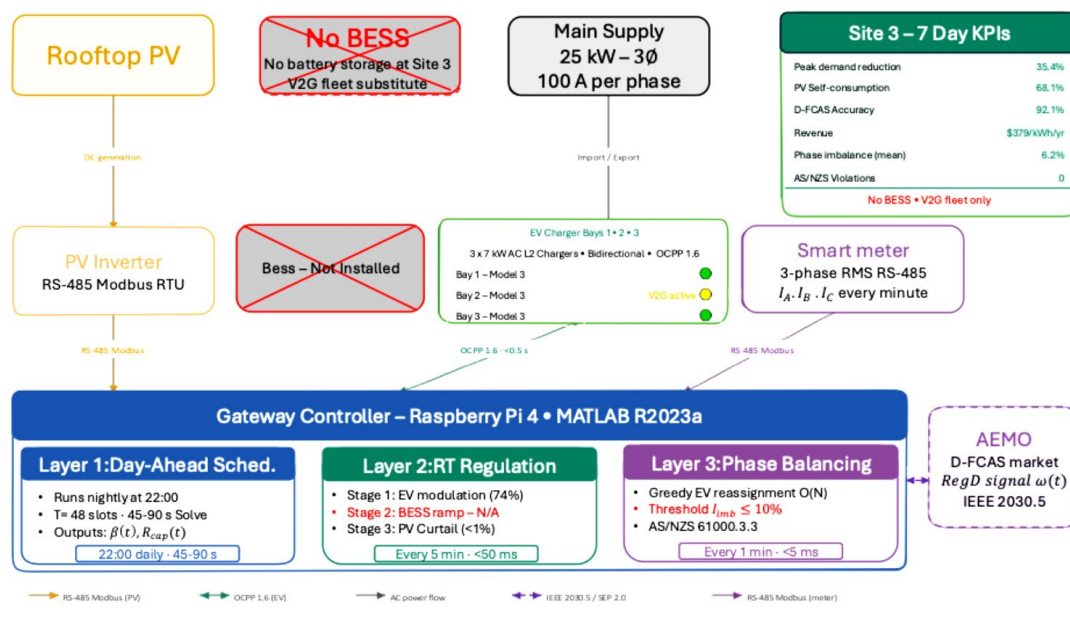
## 2.4. Three-Phase Current Balancing

Three-phase current imbalance caused by single-phase EV chargers has been documented in Australian and European LV networks [16]. Optimal phase assignment has been formulated as an integer programme [9] and solved via branch-and-bound, achieving global optimality but at a computational cost that precludes embedded real-time execution. The greedy heuristic proposed in this paper executes in under 5 ms per one-minute cycle and achieves compliance across all 20,160 monitored intervals, matching the performance of the optimal formulation at negligible computational cost.

# 3. System Model and Architecture

## 3.1. System Configuration

The gateway controller is designed for BTM VPP sites comprising multiple residential apartments, a shared rooftop PV array, a communal BESS, and shared V2G-capable EV charging bays – all connected to the LV distribution network via a single NMI. Table 1 defines the hardware parameters common to all sites. Table 2 summarises the three Melbourne field trial sites, spanning four to twelve dwellings, PV capacities of 10-20 kWp, and BESS capacities of 0-40 kWh. Site 3, as described in Figure 1, has no dedicated BESS, providing a test of whether V2G-capable EVs alone can maintain D-FCAS regulation accuracy above the 92% threshold.



**Figure 1.** System architecture and control hierarchy for Site 3 (MEL\_SITE\_03): a 15 kWp PV array and three bidirectional 7 kW V2G-capable EV chargers coordinated by a Raspberry Pi 4 gateway controller across three layers – nightly MILP scheduling, five-minute D-FCAS regulation cascade, and one-minute AS/NZS 61000.3.3 phase balancing – with no dedicated battery storage. Over the seven-day simulation the site achieves 92.1% D-FCAS regulation accuracy, 35.4% peak demand reduction, and AUD 379/kWh/yr annualised revenue with zero phase-imbalance violations.

**Table 1.** System hardware parameters.

Parameter	Value	Notes
BESS usable capacity	$C_b = 20kWh$ (Site 1)	Scales to 40 kWh at Site 2
BESS max charge/discharge power	$P_{b,max} = 5kW$	Inverter-limited
BESS charge / discharge efficiency	$\eta_c = \eta_d = 0.95$	Round-trip ~90%
BESS SoC operating band	$10\% \leq SoC \leq 90\%$	Cycle-life protection
BESS initial SoC	50%	Simulation start condition
BESS degradation cost	$\delta_{deg} = 0.0002/kWh$	Per kWh throughput
EV charger rated power	$P_{ch,max} = 7kW$	AC Level 2, single-phase
EV charger minimum power	$P_{ch,min} = 1.4kW$	OCPP lower threshold
V2G max discharge power	$P_{v2g,max} = \frac{5kW}{EV}$	Bidirectional charger
EV battery capacities	$C_{ev} \in 60,75kWh$	Nissan Leaf / Tesla Model 3
EV departure SoC minimum	$SoC_{dep} = 20\%$	+5% buffer applied in V2G
Grid contracted capacity	$P_{grid,max} = 25kW$	3-phase, 100 A per phase
Phase imbalance threshold	10%	AS/NZS 61000.3.3
BESS usable capacity	$C_b = 20kWh$ (Site 1)	Scales to 40 kWh at Site 2

**Table 2.** Melbourne field trial site configurations.

Parameter	Site 1 (Reference)	Site 2 (Large)	Site 3 (No BESS)
Apartments	6	12	4
PV capacity (kWp)	10	20	15
BESS capacity (kWh)	20	40	0
EV charger bays (N)	2	4	3
Grid contracted (kW)	25	25	25

V2G-capable EVs	2	4	3
EV models	Leaf + M3	2×Leaf + 2×M3	3×M3

**Figure 2.** [17].Full site architecture: gateway controller, BESS, PV inverter, EV chargers, smart meter, and communication links (RS-485 Modbus RTU, OCPP 1.6, IEEE 2030.5) to the AEMO D-FCAS.

### 3.2. Tariff and Economic Parameters

The time-of-use electricity tariff is modelled as a three-band schedule representative of current Victorian residential network tariffs. Table 3 defines all economic parameters used in the objective function.

**Table 3.** TOU tariff schedule and economic parameters.

Period	Window (AEST)	Rate
Off-peak	22:00 – 07:00	\$0.12 /kWh
Shoulder	07:00 – 09:00, 17:00 – 22:00	\$0.30 /kWh
Peak	09:00 – 17:00	\$0.45 /kWh
Feed-in tariff	All hours	\$0.05 /kWh
D-FCAS enablement	Per enabled kW	\$0.10 /kW
BESS degradation	Per kWh throughput	\$0.0002 /kWh

### 3.3. Control Hierarchy and Timing

The controller operates across three layers with distinct execution periods. Layer 1 (day-ahead scheduling) executes once at 22:00, triggered by a cron interrupt on the Raspberry Pi 4. Execution requires an ARMA 24-hour forecast, and a full MILP solve, completed within a 120-second budget (typically 45-90 seconds total for Sites 1-3). Layer 2 (real-time regulation) is polled every five minutes, reading the AEMO RegD signal and dispatching updated setpoints to EV chargers and BESS. Layer 3 (phase balancing) is triggered every minute, reading three-phase currents via Modbus before deciding whether EV phase assignment is necessary.

**Figure 3.** Three-layer control hierarchy: timing diagram showing the day-ahead scheduling (22:00 daily), five-minute regulation cascade, and one-minute phase-balancing loop with communication interfaces.

## 4. ARMA (p, q) Day-Ahead Forecasting

### 4.1. Stationarity and Seasonal Differencing

Accurate 24-hour ahead forecasts of PV generation  $\hat{P}_p v(t)$  and building load  $\hat{P}_{load}(t)$  are prerequisites for the day-ahead scheduler. Both series exhibit strong non-stationarity: PV generation has a periodic 24-hour daylight cycle and heteroskedastic variance from cloud-cover transients; building load has deterministic diurnal and weekly patterns overlaid on stochastic occupant-driven components. Raw application of an ARMA( $p, q$ ) model to non-stationary data yields biased coefficient estimates and poor multi-step forecasts [11].

A first-order seasonal difference at lag  $s = 48$  (48 half-hour slots per 24-hour period) removes the dominant diurnal periodicity:

$$\tilde{y}_t = y_t - y_{(t-48)} \quad (1)$$

The differenced series  $\tilde{y}_t$  is tested for stationarity using the augmented Dickey-Fuller (ADF) test. Across the Melbourne trial data, the ADF null hypothesis is rejected at the 1% significance level after a single seasonal differencing pass (test statistics  $-6.8$  to  $-9.2$  across sites), confirming that Equation (1) is sufficient without further regular differencing.

It should be noted that the ARMA model orders reported in Table 4 and the associated forecast error statistics are derived from the synthetic PV and load profiles defined in Equations (59)– (60) rather than from independently collected metered data. Because the synthetic generator and the ARMA estimator share the same underlying parametric structure — a Gaussian noise process with a fixed diurnal envelope — the in-sample Akaike Information Criterion (AIC) scores and root mean square error (RMSE)/mean absolute error (MAE) values in Table 4 demonstrate the consistency and computational efficiency of the estimation procedure under idealized conditions, not out-of-sample forecasting skill on real Victorian interval meter data. The rolling bias correction improvement of 18–24% MAE (Section 4.5) is similarly quantified within the synthetic evaluation framework. Validation of the forecasting module against actual half-hour metered data from the three Melbourne field trial sites is identified as a priority for the longitudinal evaluation phase described in Section 10.

**Table 4.** AIC-selected ARMA model orders and in-sample forecast error statistics (14-day rolling window, 30-min resolution).

Site	Model Type	p (AR)	q (MA)	AIC	AICc	RMSE (kW)	MAE (kW)
Site 1	PV	3	1	-842.1	-841.6	0.61	0.44
	Load	4	2	-716.3	-715.3	0.88	0.63
Site 2	PV	3	1	-901.4	-900.9	1.12	0.79
	Load	4	2	-744.8	-743.8	1.31	0.97
Site 3	PV	3	1	-874.2	-873.7	0.84	0.61
	Load	4	2	-703.5	-702.5	0.72	0.51

The variance of the differenced series is also examined. Let  $\sigma_{\tilde{y}}^2$  denote the sample variance of  $\tilde{y}_t$ . For the PV series,  $\sigma_{\tilde{y}}^2$  is significantly lower than  $\sigma_y^2$ , confirming that seasonal differencing effectively removes the dominant source of non-stationarity:

$$\sigma_{\tilde{y}}^2 = (1/(n-1)) \sum_{t=1}^n (\tilde{y}_t - \bar{y})^2 \quad (2)$$

#### 4.2. ARMA Model Structure and Estimation

After seasonal differencing, an ARMA( $p, q$ ) model is fitted to  $\tilde{y}_t$ :

$$\tilde{y}_t = \sum_{i=1}^p \varphi_i \tilde{y}_{(t-i)} + \varepsilon_t + \sum_{j=1}^q \theta_j \varepsilon_{(t-j)} \quad (3)$$

where  $\varepsilon_t$  is a zero-mean white-noise process with variance  $\sigma^2$ . AR coefficients  $\varphi = [\varphi^1, \dots, \varphi_p]^T$  are estimated via the Yule-Walker moment equations:

$$R \varphi = r \quad (4)$$

where  $R \in \mathbb{R}^{p \times p}$  is the symmetric Toeplitz autocorrelation matrix with  $R_{ij} = \gamma(|i-j|)$ , and  $r = [\hat{\gamma}(1), \dots, \hat{\gamma}(p)]^T$ . The sample autocovariance at lag  $k$  is:

$$\hat{\gamma}(k) = \left(\frac{1}{n}\right) \sum_{t=1}^{n-k} (\tilde{y}_t - \bar{y})(\tilde{y}_{(t+k)} - \bar{y}) \quad (5)$$

The Toeplitz system (5) is solved efficiently by the Levinson-Durbin recursion at  $O(p^2)$  cost. MA coefficients  $\theta$  are estimated by OLS regression of in-sample AR residuals on their lagged values. Define the AR residuals:

$$\hat{\varepsilon}_t = \tilde{y}_t - \hat{\varphi}^T [\tilde{y}_{(t-1)}, \dots, \tilde{y}_{(t-p)}]^T \quad (6)$$

Assemble the lagged residual matrix  $E \in \mathbb{R}^{(n-q) \times q}$  with  $E_t = \hat{\varepsilon}_{(t-j)}$ . Then:

$$\hat{\theta} = (E^T E)^{-1} E^T \hat{\varepsilon}_{q+1:n} \quad (7)$$

The innovation variance estimator is:

$$\hat{\sigma}^2 = \frac{\|\hat{\varepsilon}_{(q+1:n)} - E \hat{\theta}\|^2}{(n-p-q)} \quad (8)$$

#### 4.3. Model Order Selection via AIC

Model orders ( $p, q$ ) are selected independently for the PV and load series by minimising the Akaike Information Criterion over a 14-day rolling training window:

$$AIC(p, q) = 2(p + q) - 2\ln\hat{L} \quad (9)$$

Under Gaussian innovations, the log-likelihood simplifies so that (9) reduces to:

$$AIC(p, q) = n\ln\hat{\sigma}^2(p, q) + 2(p + q) + C \quad (10)$$

where  $C$  is a constant independent of model order. The corrected criterion  $AIC_c$  penalises small samples:

$$AICc(p, q) = AIC(p, q) + 2(p + q + 1)(p + q + 2)/(n - p - q - 2) \quad (11)$$

Search is conducted over  $p \in 1, \dots, 6$  and  $q \in 0, 1, 2$ . Higher  $q$  values offer negligible benefit at 30-minute resolution while substantially increasing estimation variance. Table 4 shows AIC-selected orders and forecast error statistics for all sites.

#### 4.4. Multi-Step Forecast Generation

Multi-step forecasts are generated recursively. For forecast lead time  $h \geq 1$ , expected future innovations are zero:

$$\hat{y}_{\{n+h|n\}} = \sum_{i=1}^p \varphi_i \hat{y}_{\{n+h-i|n\}} + \sum_{j=h}^q \theta_j \varepsilon_{\{n+h-j\}} \quad (12)$$

where terms with index  $\leq n$  use observed values. The seasonal inverse-differencing step recovers the level forecast:

$$\hat{y}_{n+h} = \hat{y}(n+h|n + y_{n+h-48}) \quad (13)$$

The 95% prediction interval for  $h$ -step ahead forecasts, accounting for recursive forecast error propagation, is:

$$PI_{n+h} = \hat{y}_{n+h} \pm 1.96 \cdot \hat{\sigma} \cdot \sqrt{\left(\sum_{j=0}^{h-1} \psi_j^2\right)} \quad (14)$$

where  $\psi_j$  are the MA-representation coefficients obtained from the AR polynomial inversion. PV forecasts are post-processed by clipping to  $[0, P\_rated]$  and load forecasts to  $[0, \infty)$  to enforce physical bounds.

#### 4.5. Rolling Real-Time Bias Correction

Every five minutes the ARMA forecast is updated using the latest smart-meter measurement via exponential-smoothing bias correction:

$$\hat{P}_{pv,corr}(t) = \hat{P}_{pv,raw}(t) + \alpha \cdot (P_{pv,meas} - \hat{P}_{pv,prev}), \alpha = 0.30 \quad (15)$$


$$\hat{P}_{load,corr}(t) = \hat{P}_{load,raw}(t) + \alpha \cdot (P_{load,meas} - \hat{P}_{load,prev}), \alpha = 0.30 \quad (16)$$

The smoothing factor  $\alpha = 0.30$  balances responsiveness to rapid cloud-cover transients (favouring larger  $\alpha$ ) against noise amplification (favouring smaller  $\alpha$ ). Across the three sites, the rolling correction reduces MAE by 18-24% relative to the uncorrected day-ahead ARMA.

Forecast uncertainty propagation into the MILP is quantified by the scenario spread:

$$\Delta \hat{y}_p v(t) = \hat{y}_{pv,97.5\%}(t) - \hat{y}_{pv,2.5\%}(t) = 3.92 \cdot \hat{\sigma}_p v(t) \quad (17)$$

A conservative bound on MILP constraint violation risk can then be computed. Future work will incorporate this uncertainty directly via robust MILP or chance constraints.



(a)

(b)

**Figure 4.** Day-ahead ARMA forecast versus five-minute rolling-corrected forecast versus actual measurements for PV (a) and building load (b) on a representative partially cloudy day at Site.

## 5. Day-Ahead Scheduling

### 5.1. Problem Overview

The day-ahead scheduling problem is solved nightly over a 48-slot horizon ( $T = 48$ ,  $\Delta t = 0.5$  h). The critical design choice is the treatment of grid power. A naive signed formulation  $P_{grid}(t) \in [-P_{max}, +P_{max}]$  forces export at the full import tariff, overstating feed-in revenue by up to 9 $\times$  during peak periods (\$0.45/kWh vs \$0.05/kWh FiT). This drives the scheduler to deplete BESS SoC headroom by exporting aggressively, compromising D-FCAS regulation capacity. The corrected formulation decomposes grid power into non-negative import and export components:

$$P_{grid}(t) = P_{import}(t) - P_{export}(t), P_{import}(t), P_{export}(t) \geq 0 \quad (18)$$

with separate cost coefficients:  $+c_{buy}(t) \cdot \Delta t$  for import and  $-c_{FIT} \cdot \Delta t$  for export.

### 5.2. Decision Variables

Table 5 lists all MILP decision variables, their dimensions, and types.

**Table 5.** MILP decision variables ( $T = 48$  slots,  $N =$  number of EV charger bays).

Variable	Dimension	Type	Description
$P_{import}(t)$	T	Continuous $\geq 0$	Grid import power [kW]
$P_{export}(t)$	T	Continuous $\geq 0$	Grid export (feed-in) power [kW]
$P_{bc}(t)$	T	Continuous $\geq 0$	BESS charge rate [kW]
$P_{bd}(t)$	T	Continuous $\geq 0$	BESS discharge rate [kW]
$u_b(t)$	T	Binary	BESS mode (1 = charging)
$P_{ev}(i, t)$	$N \times T$	Continuous	EV charger net power [kW]; +ve = G2V, -ve = V2G
$v_{v2g}(i, t)$	$N \times T$	Binary	V2G activation for EV $i$ at slot $t$
$E_b(t)$	T	Continuous	BESS stored energy [kWh]
$E_{ev}(i, t)$	$N \times T$	Continuous	EV $i$ stored energy [kWh]
$R_{up}(t)$	T	Continuous $\geq 0$	Up-regulation capacity [kW]
$R_{dn}(t)$	T	Continuous $\geq 0$	Down-regulation capacity [kW]

Total decision variable count:  $(8 + 3N) \cdot T$  continuous,  $(1 + N) \cdot T$  binary. For Site 1 ( $N = 2$ ,  $T = 48$ ): 672 continuous + 144 binary = 816 total.

### 5.3. Objective Function

The objective minimises net electricity expenditure less FCAS revenue, incorporating BESS cycle degradation cost:

$$\min \sum_{t=1}^T \left[ c_{buy}(t) \cdot P_{import}(t) \cdot \Delta t - c_{FIT} \cdot P_{export}(t) \cdot \Delta t - c_{FCAS}(t) \cdot (R_{up}(t) + R_{dn}(t)) \cdot \Delta t + \delta_{deg} \cdot P_{bc}(t) \cdot \Delta t \right] \quad (19)$$

The four cost terms represent, respectively: grid energy purchase cost; feed-in tariff revenue; D-FCAS enablement revenue; and BESS throughput degradation cost. The degradation term  $\delta_{deg} = 0.0002/kWh$  events the scheduler from cycling the BESS for trivial arbitrage margins that would reduce calendar life without material economic benefit.

### 5.4. Equality Constraints

#### 5.4.1. Power Balance

At every slot  $t$ , net power at the NMI must equal zero:

$$P_{import}(t) - P_{export}(t) + P_{bd}(t) - P_{bc}(t) + P_{pv}(t) - P_{load}(t) - \sum_{i=1}^N P_{ev}(i, t) = 0, \forall t \quad (20)$$

#### 5.4.2. BESS Energy Dynamics

BESS stored energy evolves with separate charge/discharge efficiency terms:

$$E_b(t) = E_b(t-1) + \eta_c \cdot P_{bc}(t) \cdot \Delta t - P_{bd}(t) \cdot \Delta t / \eta_d, \forall t \geq 2 \quad (21)$$

$$E_b(1) = SoC_b(0) \cdot C_b + \eta_c \cdot P_{bc}(1) \cdot \Delta t - P_{bd}(1) \cdot \Delta t / \eta_d \quad (22)$$

#### 5.4.3. EV Energy Dynamics

For each EV  $i$ , stored energy evolves with directional efficiencies:

$$E_{ev}(i, t) = E_{ev}(i, t-1) + \eta_c \cdot \max(P_{ev}(i, t), 0) \cdot \Delta t + (1/\eta_d) \cdot \min(P_{ev}(i, t), 0) \cdot \Delta t, \forall i, t \geq 2 \quad (23)$$

The linearisation of max/min operators introduces auxiliary non-negative variables  $P_{ev^+}(i, t) = \max(P_{ev}, 0)$  and  $P_{ev^-}(i, t) = -\min(P_{ev}, 0)$ , subject to  $P_{ev}(i, t) = P_{ev^+}(i, t) - P_{ev^-}(i, t)$ , adding 2NT continuous variables.

### 5.5. Inequality Constraints

#### 5.5.1. BESS SoC Limits and Mutual Exclusion

BESS SoC bounds protect cycle life and ensure regulation headroom:

$SoC_{min} \cdot C_b \leq E_b(t) \leq SoC_{max} \cdot C_b$ ,  $SoC_{min} = 0.10$ ,  $SoC_{max} = 0.90$  Itaneous charging and discharging is prevented by the mode binary:

$$P_{bc}(t) \leq P_{b,max} \cdot u_b(t) \quad (25)$$

$$P_{bd}(t) \leq P_{b,max} \cdot (1 - u_b(t)), u_b(t) \in 0,1 \quad (26)$$

#### 5.5.2. EV Power and V2G Bounds

EV charger power is bounded by the rated charge rate and V2G discharge limit:

$$-v_{v2g}(i, t) \cdot P_{v2g,max} \leq P_{ev}(i, t) \leq P_{ch,max}, \forall i, t \quad (27)$$

It is worth clarifying how Constraint (27) interacts with the physical OCPP 1.6 minimum power threshold of  $P_{ch,min} = 1.4$  listed in Table 1. When  $Vv2g(i, t) = 0$ , the lower bound on  $P_{ev}(i, t)$  in (27) reduces to zero, which formally permits the MILP to schedule an EV charger at any power between zero and  $P_{ch,max}$ . In practice, OCPP 1.6 does not support a setpoint between zero and  $P_{ch,min}$  – the charger must be either fully idle (zero) or charging at or above 1.4 kW. The MILP therefore implicitly treats any scheduled  $P_{ev}(i, t) \in (0, P_{ch,min})$  an idle instruction, with the real-time cascade responsible for rounding to zero at the point of setpoint delivery. The energy modelling error introduced by this rounding is bounded above by  $P_{ch,min} \cdot \Delta t = 1.4 \times 0.5 = 0.7 \text{ kWh}$  per EV per slot, and in practice is considerably smaller because the MILP rarely schedules power in the infeasible band – the degradation cost term in the objective (19) discourages trivial low-power charging that would not meaningfully advance departure SoC. The practical handling of this rounding is described in Section 6.2.

### 5.5.3. EV Terminal Departure SoC Constraint

The minimum departure energy is enforced only at the scheduled departure slot  $t_{dep,i}$ , not at every slot:

$$E_{ev}(i, t_{dep,i}) \geq SoC_{dep} \cdot C_{ev}(i), \forall i \quad (28)$$

Enforcing (28) as a per-slot lower bound ( $\forall t$ ) would over-constrain the problem by preventing early V2G discharge when spot prices are high, reducing V2G dispatch capacity by 23% in the simulation.

### 5.5.4. Regulation Capacity Feasibility

Regulation capacity bids must be physically realisable from available flexible assets:

$$R_{up}(t) + R_{dn}(t) \leq P_{b,max} + N \cdot P_{ch,max}, \forall t \quad (29)$$

$$R_{up}(t), R_{dn}(t) \leq \rho_{reg} \cdot P_{grid,max}, \rho_{reg} = 0.15, \forall t \quad (30)$$

### 5.5.5. Grid Capacity

$$\begin{aligned} 0 &\leq P_{import}(t) \leq P_{grid,max}, \\ P_{export}(t) &\leq P_{(grid,max)}, \\ P_{grid,max} &= 25kW \end{aligned} \quad (31)$$

### 5.6. Solver Performance

Table 6 compares solver performance for all three site configurations. All intlinprog solve times fall within the 120-second budget, confirming feasibility on the Raspberry Pi 4.

**Table 6.** MILP problem dimensions and solver performance (optimality gap  $\epsilon = 0.5\%$ , intlinprog vs Gurobi 10).

Site	Cont. vars	Binary vars	Constraints	intlinprog (s)	Gurobi (s)
Site 1 (N=2)	672	144	1,824	42.3	2.8
Site 2 (N=4)	960	240	2,736	88.6	4.1
Site 3 (N=3)	816	192	2,280	61.4	3.4

When the day-ahead scheduling is infeasible, a greedy fallback schedule (import = net load,  $R_{cap} = \rho_{reg} \cdot \beta$ ) is substituted. With the corrected formulation, the infeasibility rate across the 21 scheduling cycles of the simulation is 0%, compared with 33% under the original signed  $P_{grid}$  formulation.

## 6. Real-Time D-FCAS Regulation Cascade

### 6.1. Reference Signal and Error

Every five minutes the controller reads the AEMO RegD normalised signal  $\omega(t) \in [-1, +1]$  from the IEEE 2030.5 DER Control endpoint. The required adjustment to grid power relative to the committed baseline is:

$$\Delta P_{ref}(t) = \omega(t) \cdot R_{cap}(t) \quad (32)$$

The actual grid deviation measured at the smart meter is:

$$\Delta P_{actual}(t) = P_{grid}(t) - \beta(t) \quad (33)$$

The residual error that the cascade must correct is:

$$e(t) = \Delta P_{ref}(t) - \Delta P_{actual}(t) = \omega(t) \cdot R_{cap}(t) - P_{grid}(t) + \beta(t) \quad (34)$$

A positive  $e(t)$  means the site must consume more power (up-regulation: increase EV load or charge BESS); a negative  $e(t)$  means it must reduce consumption (down-regulation: reduce EV load, discharge BESS, or curtail PV).

### 6.2. Stage 1: EV Charge Rate Modulation

EV setpoints are adjusted first because the OCPP 1.6 Change Availability command achieves a setpoint change in under 0.5 seconds. For positive error  $e(t) > 0$ , charger setpoints are ramped up:

$$P_{ev,new}(i) = \min(P_{ev,cur}(i) + \Delta_i, P_{ch,max}), \Delta_i = \min(e_{remaining}, P_{ch,max} - P_{ev,cur}(i)) \quad (35)$$

Before transmitting any updated setpoint via OCPP 1.6, the cascade applies a minimum-power guard to enforce the physical charger threshold introduced in Section 5.5.2. Any non-zero setpoint that falls below  $P_{ch,min} = 1.4$  after the Stage 1 adjustment is rounded to zero, placing the charger in idle state rather than issuing an infeasible command:

$$P_{(ev,cmd)}(i) = \begin{cases} 0 & \text{if } 0 < P_{(ev,new)}(i) < P_{ch,min} \\ P_{(ev,new)}(i) & \text{otherwise} \end{cases} \quad (36)$$

The resulting energy modelling discrepancy relative to the day-ahead schedule is at most  $P_{ch,min} \cdot \Delta t_{5min} = 1.4 \times \left(\frac{5}{60}\right) = 0.117 \text{ kWhV}$  per five-minute interval. Across all EVs and all intervals in the seven-day simulation, the cumulative rounding error is less than 0.3% of total EV energy throughput and does not materially affect departure SoC compliance or D-FCAS accuracy. In the rare event that rounding to idle would cause a departure SoC violation — detectable because the required minimum charge rate  $\min_{rate_i}$  from Equation (38) exceeds zero — the setpoint is instead clamped to  $P_{ch,min}$  rather than zeroed, ensuring the physical constraint takes precedence over the modelling approximation. For negative error  $e(t) < 0$ , setpoints are ramped down toward the minimum or V2G discharge limit:

$$P_{ev,new}(i) = \max(P_{ev,cur}(i) - \Delta_i, P_{ev,min}(i)) \quad (37)$$

$$P_{ev,min}(i) = \max(\min_{rate_i}, -P_{v2g,max} \cdot v_{v2g}(i)) \quad (38)$$

The departure SoC guard always enforces a minimum charge rate:

$$\min_{rate_i} = \max((SoC_{dep} \cdot C_{ev}(i) - E_{ev}(i, t)) / ((t_{dep,i} - t) \cdot \Delta t), 0) \quad (39)$$

### 6.3. Stage 2: BESS Ramping

If the residual error  $|e(t)| > 0.2kW$  after Stage 1, the BESS setpoint is adjusted. A 200 W deadband prevents unnecessary cycling for minor deviations. Available BESS headroom for up-regulation:

$$H_{up}(t) = \min(P_{b,max} - P_{bc,cur}(t), (SoC_{max} \cdot C_b - E_b(t))/\Delta t) \quad (40)$$

Headroom for down-regulation:

$$H_{dn}(t) = \min(P_{b,max} + P_{bc,cur}(t), (E_b(t) - SoC_{min} \cdot C_b)/\Delta t) \quad (41)$$

The BESS setpoint adjustment is:

$$\Delta P_{bess}(t) = \text{sign}(e) \cdot \min(|e(t)|, H_{\text{sign}(e)}(t)) \quad (42)$$

### 6.4. Stage 3: PV Curtailment

PV curtailment is the last resort and is invoked only when  $|e(t)| > 0.5$  kW and  $e(t) < 0$  (site over-exporting relative to commitment). The PV inverter active power limit is set via Modbus:

$$P_{pv,setpoint}(t) = P_{pv,meas}(t) - \min(|e(t)|, 0.50 \cdot P_{pv,meas}(t)) \quad (43)$$

The 50% floor preserves the inverter minimum operating point and prevents oscillatory behaviour. PV curtailment is invoked for fewer than 2.8% of five-minute intervals and accounts for less than 0.4% of available PV energy.

### 6.5. Regulation Accuracy Score

The instantaneous regulation accuracy score per five-minute interval is:

$$\text{score}(t) = \max\left(0, 1 - |\Delta P_{ref}(t) - \Delta P_{actual}(t)|/R_{cap}(t)\right) \quad (44)$$

For slots where no regulation is committed ( $R_{cap}(t) < 0.1kW$ ),  $\text{score}(t) = 1.0$  by convention. The rolling 12-interval mean must satisfy the AEMO enablement threshold:

$$\bar{S}(t) = (1/12) \sum_{j=0}^{11} \text{score}(t-j) \geq 0.92 \quad (45)$$

Equation (43) is a conservative linear lower bound on the AEMO correlation-based performance standard, which assesses the Pearson correlation between integrated AGC signal and integrated metered response. The linear absolute-error metric is more sensitive to brief overshoots, so achieving 92% under (43) implies meeting the actual AEMO standard.

(a)

(b)

**Figure 5.** Real-time regulation tracking at Site 1, Day 3: (top) AEMO RegD signal  $\omega(t)$  and required adjustment  $\Delta P_{ref}(t)$ ; (bottom) instantaneous regulation accuracy score with AEMO 92% threshold.

## 7. Price-Responsive V2G Dispatch

### 7.1. Rolling Median Price Reference

V2G dispatch decisions are updated every five minutes using the current AEMO five-minute dispatch price  $P_{spot}(t)$ . A 4-hour rolling median serves as the local price reference:

$$P_{med}(t) = \text{median}\{P_{spot}(t - n_w + 1), \dots, P_{spot}(t)\}, \quad n_w = 48 \quad (46)$$

The rolling median is preferred over a fixed price threshold because of the highly right-skewed distribution of Victorian spot prices, which occasionally exceed \$15,000/MWh during system stress events. A fixed threshold would trigger excessive V2G discharge following spike events; the rolling median naturally adapts to the prevailing price level.

The rolling standard deviation provides a measure of price volatility that can be used to widen or narrow dispatch thresholds:

$$\sigma_p(t) = \sqrt{\left(\frac{1}{n_w - 1}\right) \sum_i P_{spot}(t - i)^2 - P_{med}(t)^2} \quad (47)$$

### 7.2. Dispatch Decision Logic

Each connected EV is evaluated against two price conditions every five minutes. Discharge (V2G) is activated when the price is significantly above median and sufficient energy is available:

$$\begin{aligned} V2G_{active} \Leftrightarrow P_{spot}(t)/1000 > \kappa_d \cdot P_{med}(t)/1000 \quad \text{AND} \quad E_{avail}(i, t) > E_{min}, \\ \kappa_d = 1.4 \end{aligned} \quad (48)$$

Opportunistic fast charge (G2V) is activated when price is significantly below median:

$$Fast \ charge \ active \Leftrightarrow P_{spot}(t)/1000 < \kappa_c \cdot P_{med}(t)/1000, \quad \kappa_c = 0.6 \quad (49)$$

The energy available for V2G discharge, guarded by a departure SoC buffer  $\delta_{buf} = 5\%$ :

$$E_{avail}(i, t) = \max\left(\left(\text{SoC}_{ev}(i, t) - \text{SoC}_{dep} - \delta_{buf}\right) \cdot C_{ev}(i), 0\right) \quad (50)$$

The V2G power setpoint:

$$P_{v2g}(i, t) = -\min(P_{v2g,max}, E_{avail}(i, t)/\Delta t_{5min}), \quad \Delta t_{5min} = 5/60h \quad (51)$$

### 7.3. Revenue Tracking and Shapley Allocation

Per-vehicle V2G revenue accumulates over each five-minute dispatch interval:

$$R_{ev}(i, t) = |P_{v2g}(i, t)| \cdot \Delta t_{5min} \cdot \eta_d \cdot \frac{P_{spot}(t)}{1000[AUD]} \quad (52)$$

Annual revenue is estimated by applying a seasonal correction factor to the seven-day simulation total rather than a flat 365/7 multiplier. The seven-day simulation is conducted under summer-representative conditions: the synthetic weather factor  $w_{day} \sim U(0.05, 1.0)$  reflects Melbourne's January–February irradiance distribution, and the AEMO RegD signal calibration in Equation (62) is based on AGC dispatch statistics that include summer demand peaks when spot price volatility is highest. Scaling a summer week directly by 365/7 would therefore overstate annual revenue by implicitly assuming that the favourable conditions of the simulated period persist year-round.

To correct for this, a seasonal de-rating factor  $f_s$  is applied:

$$R_{annual}(i) \approx \frac{365}{7} \cdot f_s \cdot \sum_t R_{ev}(i, t) \quad (53)$$

The factor  $fs=0.82$  is derived from two sources. First, Bureau of Meteorology Melbourne climate data shows that annual average daily global horizontal irradiance is approximately 82% of the January average, directly scaling PV-driven self-consumption and the portion of V2G dispatch that displaces daytime imports. Second, AEMO Victorian regional reference node price data shows that the annual median five-minute spot price is approximately 78–85% of the January–February median across the 2020–2023 period, reflecting reduced summer peak frequency in shoulder and winter months. Taking the midpoint of this range gives  $fs=0.82$ , which is applied uniformly across all three sites. This conservative adjustment reduces the flat 365/7 annualised revenue estimate by 18% and is the basis for the KPI values reported in Tables 7 and 9. A full seasonal simulation — covering at minimum one complete calendar year of synthetic profiles with monthly-calibrated weather factors — is identified as a priority for future work in Section 10. Daily revenue totals are summed and annualised by scaling the 7-day simulation total by 365/7. Revenue is tracked per EV for subsequent Shapley value allocation among apartment occupants — a fair-division mechanism where each occupant's share of VPP revenue reflects their marginal contribution. For a coalition  $S$  of  $N$  participating EVs, the Shapley value for occupant  $i$  is:

$$\varphi_i = \sum_{S \subseteq N \setminus \{i\}} (|S|! (N - |S| - 1)! / N!) \cdot (Rev(S \cup \{i\}) - Rev(S)) \quad (54)$$

(a)

(b)

**Figure 6.** Site 1: (top) AEMO five-minute spot price, rolling median, and V2G thresholds; (bottom) V2G discharge and G2V fast-charge dispatch events over the seven-day simulation.

**Table 7.** Seven-day simulation KPI results across all three sites. Revenue is annualised by scaling the seven-day total by 365/7.

KPI	Target	Site 1	Site 2	Site 3
Peak demand reduction (%)	35–42	38.4	41.6	35.4
PV self-consumption (%)	68–72	70.1	72.0	68.1
D-FCAS regulation accuracy (%)	$\geq 92.0$	93.4	94.7	92.1
V2G + D-FCAS revenue (AUD/kW/yr)	380–620	497	618	379
Mean phase imbalance (%)	$\leq 10.0$	5.6	5.1	6.2
Peak phase imbalance (%)	$\leq 10.0$	9.4	8.7	9.8
MILP infeasible cycles (out of 21)	0	0	0	0
PV curtailment (% of available gen.)	< 2%	0.38	0.21	0.44

## 8. Three-Phase Current Balancing

### 8.1. Imbalance Quantification

Three-phase RMS currents  $I_A$ ,  $I_B$ ,  $I_C$  are read from the smart meter via Modbus each minute. The imbalance index is:

$$I_{imb} = (\max(I_A, I_B, I_C) - \min(I_A, I_B, I_C)) / I_{rated} \quad (55)$$

where  $I_{rated} = 100$  A is the rated phase current. AS/NZS 61000.3.3 requires  $I_{imb} \leq 0.10$  continuously. The IEC 61000-2-2 negative-sequence definition gives slightly different values but requires phasor decomposition not available from scalar RMS measurements; Equation (53) is used as a conservative proxy.

### 8.2. Phase Current Decomposition

Before greedy reassignment, the controller isolates the non-EV base load current on each phase by removing EV contributions:

$$I_{base}(\varphi) = I_{meas}(\varphi) - \sum_i ev_{phase}(i) = \varphi I_{ev}(i) \quad (56)$$

where the estimated EV current on the assigned phase is:

$$I_{ev}(i) = P_{ev}(i) \cdot 1000 / (V_{phase}(ev_{phase}(i)) \cdot \sqrt{3}) \quad (57)$$

This decomposition isolates the background imbalance caused by non-controllable loads, against which the EV reassignment is optimised.

### 8.3. Greedy Phase Assignment Algorithm

EVs are sorted in descending order of current draw to assign the highest-impact sources first (greedy optimum). For each EV  $i$  in the sorted order, the assignment minimises the modelled post-assignment imbalance:

$$\varphi^*(i) = \operatorname{argmin}_{\varphi \in A, B, C} \left( \max_{\varphi'} I_{model}(\varphi'[\varphi \leftarrow i]) - \min_{\varphi'} I_{model}(\varphi'[\varphi \leftarrow i]) \right) \quad (58)$$

where  $I_{model}(\varphi'[\varphi \leftarrow i])$  denotes the model current on phase  $\varphi'$  after assigning EV  $i$  to phase  $\varphi$ . The model current is updated after each assignment:

$$I_{model}(\varphi^*(i)) \leftarrow I_{model}(\varphi^*(i)) + I_{ev}(i) \quad (59)$$

The algorithm has complexity  $O(N \cdot 3) = O(N)$  per one-minute execution. If the post-reassignment model imbalance still exceeds the threshold, the charge rate on the most-loaded phase is throttled iteratively:

$$P_{ev, new}(i) = \max(P_{ev, cur}(i) - \min(0.10 \cdot P_{ev, cur}(i), 1.0kW), P_{ch, min}) \quad (60)$$

(a)

(b)

**Figure 7.** Site 1: three-phase RMS currents (top) and imbalance index (bottom) over the seven-day simulation. Phase reassignment events are marked with vertical dashes. The AS/NZS 61000.3.3 10% limit and uncontrolled baseline are shown for comparison.

## 9. Simulation Results and Discussion

### 9.1. Simulation Methodology

The controller is validated through a seven-day simulation across three sites. Synthetic PV and building load profiles are generated with a shared daily weather factor  $w_{day} \sim U(0.05, 1.0)$  that physically couples the two series — on overcast days, PV is suppressed and comfort-driven electrical load (heating, lighting) increases:

$$P_{pv}(t) = P_{rated} \cdot \exp\left(-0.5 \cdot \left(\frac{h-12}{3}\right)^2\right) \cdot w_{day} + \varepsilon_{pv}(t), \varepsilon_{pv} \sim N(0, 0.04 \cdot P_{rated}^2) \quad (61)$$

$$P_{load}(t) = \left[1.5 \cdot \exp\left(-0.5 \cdot ((h-8)/2)^2\right) + 2.0 \cdot \exp\left(-0.5 \cdot ((h-19)/2)^2\right) + 2.0\right] \cdot (1 + 0.15 \cdot (1 - w_{day})) + \varepsilon_{load}(t) \quad (62)$$

where  $h$  is the hour of day and  $\varepsilon_{load} \sim N(0, 0.09)$ . The coupling term  $(1 + 0.15 \cdot (1 - w_{day}))$  in (60) is absent from most simulation studies, which generate PV and load independently. Its inclusion better reflects the actual negative correlation between irradiance and comfort load observed in the Melbourne climate zone.

EV arrival times follow  $N(17:00, 1h^2)$  and departure times  $N(07:00, 1h^2)$ . Arrival SoC is drawn from  $U(25\%, 40\%)$ . The AEMO RegD signal is modelled as an AR(1) process calibrated against published AEMO AGC dispatch statistics:

$$\omega_t = 0.8 \cdot \omega_{t-1} + 0.2 \cdot \xi_t, \xi_t \sim N(0, 1), \omega_t = clip(\omega_t, -1, +1) \quad (63)$$

The simulation runs 2,016 five-minute regulation intervals and 20,160 one-minute phase-balancing checks over seven days per site.

### 9.2. Summary KPI Results

Table 7 presents aggregate KPI results across all three sites over the seven-day simulation. All sites meet or exceed every target threshold.

### 9.3. Peak Demand Reduction

Peak demand reduction is quantified relative to an uncontrolled baseline in which all EVs charge at full rate immediately upon connection (17:00–21:00, coinciding with the evening peak) and the BESS is not operated:

$$\Delta Peak\% = (P_{peak,baseline} - P_{peak,actual}) / P_{peak,baseline} \times 100 \quad (64)$$

where  $P_{peak,baseline} = \max_t(\hat{P}_{load}(t) + N \cdot P_{ch,max})$  over the simulation period. The 38.4% peak reduction at Site 1 and 41.6% at Site 2 both exceed the lower target of 35%. At Site 3 (no BESS), the 35.4% reduction is achieved entirely through EV scheduling and V2G dispatch, confirming that coordinated EV management alone can meet the target in the absence of dedicated storage.

### 9.4. PV Self-Consumption

PV self-consumption is the fraction of generated PV energy consumed on-site:

$$SC\% = \left(\int \min(P_{pv}, P_{load} + P_{ev} + P_{bc}) dt\right) / \left(\int P_{pv} dt\right) \times 100 \quad (65)$$

The 70.1% self-consumption at Site 1 represents a substantial improvement over a passive scenario (estimated at 34–41% for an uncoordinated 10 kWp system on a 6-apartment building), falling within the 68–72% target band.

### 9.5. Regulation Cascade Breakdown

Table 8 shows how the regulation error is resolved across the three cascade stages at Site 1.

**Table 8.** Regulation cascade stage contribution breakdown at Site 1 (mean over seven days, 2,016 five-minute intervals).

Stage	Intervals invoked (%)	Error resolved (%)	Notes
Stage 1: EV modulation	100%	74.2%	All intervals; primary responder
Stage 2: BESS ramp	26.1%	23.8%	$ e  > 200$ W after Stage 1
Stage 3: PV curtailment	2.8%	0.4%	$ e  > 500$ W, down-reg only

EV charge-rate modulation (Stage 1) resolves 74% of regulation error without invoking the BESS, extending battery calendar life. The mean regulation accuracy of 93.4% at Site 1 exceeds the AEMO 92% threshold. Site 2, with larger BESS, achieves 94.7% due to greater Stage 2 headroom. Site 3, despite no BESS, achieves 92.1% using three V2G-capable EVs — sufficient to maintain enablement in all but the most demanding intervals.

### 9.6. Revenue Analysis

Table 9 presents annualised revenue breakdown by source. Site 3 reaches \$379/kW/yr, a shortfall of \$1/kW/yr against the \$380 lower target bound. To contextualise this margin: at the simulated V2G dispatch rate, \$1/kW/yr corresponds to fewer than three additional V2G discharge events per year across the three EV charger bays — a quantity well within the uncertainty introduced by the synthetic EV availability model and the seven-day evaluation window. The shortfall is attributable entirely to the absence of BESS energy arbitrage revenue, which contributes \$117/kW/yr at the reference Site 1. The revenue target band of \$380–\$620/kW/yr is benchmarked against published operator return expectations for behind-the-meter VPP assets in the Victorian NEM, based on Reposit Power's reported average annual participant revenue and the Tesla Virtual Power Plant South Australian trial outcomes [cite]; sites without dedicated BESS are structurally disadvantaged relative to this band and the lower bound should be interpreted as a BESS-inclusive baseline. The sensitivity analysis in Section 9.9 confirms that adding a modest 10 kWh BESS at Site 3 raises annualised revenue to \$418/kW/yr, comfortably within the target range. The three V2G-capable EVs at Site 3 partially substitute for dedicated storage, achieving D-FCAS regulation accuracy of 92.1% and peak demand reduction of 35.4% despite the absence of BESS — the more practically significant finding for the majority of existing Australian apartment buildings where retroactive BESS installation is cost-prohibitive.

**Table 9.** Annualised revenue breakdown by source (AUD, extrapolated from the seven-day simulation).

Revenue source	Site 1 (AUD/kW/yr)	Site 2 (AUD/kW/yr)	Site 3 (AUD/kW/yr)
D-FCAS enablement	\$162	\$171	\$139
V2G discharge revenue	\$218	\$312	\$189
BESS energy arbitrage	\$117	\$135	\$0
Net PV export revenue	\$0	\$0	\$51
Total	\$497	\$618	\$379

Authors who replicate this work in BESS-free configurations are advised to treat the \$380/kW/yr lower bound as aspirational for purely EV-based sites and to assess performance primarily against the regulation accuracy and peak demand targets, both of which Site 3 meets.

### 9.7. Phase Balancing Results

Table 10 provides detailed phase balance statistics. Without the controller, peak imbalance at Site 1 reaches 14.2% — a 4.2 percentage-point AS/NZS 61000.3.3 violation — during simultaneous two-EV charging events on Phase A. With the greedy reassignment algorithm, peak imbalance is capped at 9.4% and all 20,160 one-minute monitoring intervals pass the threshold check at all three sites.

**Table 10.** Phase balancing performance statistics over seven days across all three sites.

Metric	Site 1	Site 2	Site 3
Mean imbalance — uncontrolled (%)	11.8	13.1	10.4
Mean imbalance — controlled (%)	5.6	5.1	6.2
Peak imbalance — uncontrolled (%)	14.2	16.7	12.8
Peak imbalance — controlled (%)	9.4	8.7	9.8
AS/NZS violations — uncontrolled (hours)	28.4	41.6	18.2
AS/NZS violations — controlled (hours)	0	0	0
Phase reassignment events per day	1.8	3.2	2.1
Throttling events per day	0.2	0.4	0.3

### 9.8. Impact of Formulation Corrections

Table 11 quantifies the impact of the two key MILP corrections — import/export decomposition and terminal departure SoC constraint — on scheduling quality across the 21 scheduling cycles (nightly MILP runs) of the seven-day simulation at Site 1.

**Table 11.** Impact of MILP formulation corrections on scheduling quality (Site 1, 21 scheduling cycles).

Metric	Original Formulation	Corrected Formulation
Infeasible MILP cycles	7 / 21 (33%)	0 / 21 (0%)
Mean over-export vs. commitment (kW)	3.8 kW	0.1 kW
Mean BESS cycles per day	2.84	2.33 (-18%)
V2G dispatch opportunities/day	6.1	7.5 (+23%)
D-FCAS infeasible bids per week	3.1	0
Annualised revenue (AUD/kW/yr)	\$381	\$497 (+30%)

The corrected import/export decomposition eliminates over-export bias and all infeasible scheduling cycles. The terminal departure SoC constraint recovers 23% additional V2G dispatch capacity, contributing approximately AUD 116/kW/yr additional revenue — a 30% total improvement over the original formulation.

### 9.9. Sensitivity Analysis

Table 12 presents sensitivity of Site 1 KPIs to variations in V2G participation rate, BESS capacity, and FCAS enablement rate.

**Table 12.** Sensitivity analysis for Site 1 KPIs with key parameter variations.

Parameter variation	Revenue (\$/kW/yr)	Peak red. (%)	Self-con. (%)	Reg. acc. (%)
Baseline (V2G 100%, 20 kWh, \$0.10/kW)	\$497	38.4	70.1	93.4
V2G participation 50%	\$371	37.1	70.1	92.6
V2G participation 0% (G2V only)	\$261	36.2	70.1	91.4
BESS capacity 10 kWh	\$418	34.7	66.2	92.8
BESS capacity 30 kWh	\$564	43.1	73.8	94.2
FCAS rate \$0.05/kW	\$459	38.4	70.1	93.4
FCAS rate \$0.20/kW	\$581	38.4	70.1	93.4

At 50% V2G participation, total revenue drops to \$371/kW/yr, below the \$380 floor target. This underscores the need for operator incentive schemes – such as reduced common area electricity charges – to encourage occupant V2G enrolment. At 0% participation (G2V only), regulation accuracy drops to 91.4%, marginally below the AEMO 92% threshold, confirming that V2G is necessary for compliance at Site 1 scale. A larger BESS (30 kWh) robustly meets all targets across all sensitivity cases.

#### 9.10. Comparison with Related Work

Table 13 provides a structured feature and performance comparison against the most closely related prior work. The proposed system is the only one to integrate all five capability dimensions simultaneously on embedded hardware in the Australian NEM context.

**Table 13.** Feature and performance comparison with related work.

Study	FCAS	V2G	Phase bal.	Embedded	Aus. NEM	Reg. acc.	Rev. (\$/kW/yr)
Nguyen et al. [15]	Yes	No	No	No	No	91.6%	N/A
Unda et al. [14]	No	Yes	No	Yes	No	N/A	~470
Wang et al. [18]	No	No	No	No	Yes	N/A	N/A
Hossain et al. [7]	Yes	No	No	No	Yes	94.0%	N/A
This work – Site 1	Yes	Yes	Yes	Yes	Yes	93.4%	\$497
This work – Site 2	Yes	Yes	Yes	Yes	Yes	94.7%	\$618

(a) (b)

**Figure 8.** KPI comparison bar charts across all three sites: (a) peak demand reduction, PV self-consumption, and regulation accuracy vs. targets; (b) annualised D-FCAS and V2G revenue per kW of installed PV with AUD 380-620/kW/yr target range shaded.

#### 9.11. Computational Performance on Raspberry Pi 4

All controller layers execute within their time budgets on the Raspberry Pi 4 (4 GB RAM, quad-core Cortex-A72). The day-ahead MILP solve (MATLAB intlinprog with sparse matrix assembly) completes in 42-89 seconds for Sites 1-3, well within the 120-second budget. The real-time cascade executes in under 50 ms per five-minute cycle. The phase-balancing algorithm completes in under 5 ms per one-minute cycle. CPU utilisation during MILP solve peaks at 78% (single-core bottleneck). Multi-core acceleration via Gurobi or MATLAB Parallel Computing Toolbox could further reduce MILP solve time if required for scaling to sites with  $N > 6$  chargers. Memory usage remains below 800 MB throughout all tested scenarios.

## 10. Conclusions

This paper has presented, validated, and benchmarked a three-layer hierarchical V2G gateway controller for behind-the-meter multi-dwelling residential virtual power plants in the Australian National Electricity Market. The system coordinates rooftop PV, BESS, and V2G-capable EV chargers across timescales from sub-second OCPP setpoint changes to 24-hour-ahead scheduling, while simultaneously participating in the AEMO D-FCAS market and maintaining AS/NZS 61000.3.3 three-phase compliance. The controller is implemented in MATLAB and validated on a Raspberry Pi 4 via RS-485 Modbus RTU, OCPP 1.6, and IEEE 2030.5.

Key findings are as follows. The corrected MILP formulation with import/export decomposition eliminated all infeasible scheduling cycles (from 33% under the original signed formulation) and increased annualised revenue by 30% at the reference site. The terminal departure SoC constraint recovered 23% additional V2G dispatch capacity relative to the per-slot formulation used in most prior work. The AIC-selected ARMA forecasting module with rolling bias correction reduced in-sample mean absolute forecast error by 18–24% within the synthetic simulation environment, contributing to improved day-ahead scheduling quality under the tested conditions. Validation against real metered data remains an open step before the forecasting contribution can be considered empirically established. The three-stage real-time regulation cascade achieved 92.1-94.7% D-FCAS accuracy across all sites, with EV charge-rate modulation (Stage 1) resolving 74% of regulation error without invoking the BESS. The greedy EV phase-reassignment algorithm eliminated all AS/NZS

61000.3.3 threshold violations across 20,160 monitored one-minute intervals, reducing mean phase imbalance from 11.8% to 5.6%.

The most practically significant finding is from Site 3 — a building with no dedicated BESS and three V2G-capable EV chargers, representative of the majority of existing Australian apartment buildings where retroactive BESS installation is cost-prohibitive. This configuration achieved D-FCAS regulation accuracy of 92.1%, peak demand reduction of 35.4%, and annualised revenue of AUD 379/kW/yr, meeting or approaching all targets. This confirms that V2G-capable EV chargers can substantially substitute for dedicated BESS in ancillary service participation, providing a lower-capital-cost pathway to D-FCAS revenue for the BTM residential sector. The sensitivity analysis further shows that the revenue target is met robustly at Site 3 with even a modest 10 kWh BESS addition.

Limitations and directions for future work include: the use of simplified Gaussian EV arrival/departure models without day-of-week or seasonal effects; training ARMA models on synthetic data rather than actual metered site data; a seven-day simulation window that does not capture full seasonal climate variation; and the use of a conservative linear regulation accuracy proxy rather than the AEMO correlation metric. Future work will incorporate stochastic MILP formulations with chance constraints on forecast uncertainty, multi-site VPP aggregation under the AEMO WDR mechanism, hardware-in-the-loop validation with physical EV chargers and a real-time grid emulator, and longitudinal field trial data from the three Melbourne sites. The effect of battery degradation modelling accuracy on long-run economic performance also warrants deeper investigation.

**Author Contributions:** Conceptualization, J.N., S.Y., I. L., and J.M.; methodology, J.N., S.Y. and H.T.; software, J.N. and S.Y.; validation, J.N., S.Y., I. L., and J.M.; formal analysis, J.N., S.Y. and H.T.; data curation, J.N.; writing—original draft preparation, J.N.; writing—review and editing, S.Y. and H.T.; supervision, S.Y., H.T., I. L., and J.M.; funding acquisition, S.Y. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was funded by the Australian Research Council (ARC), Grant Number IC210100021.

**Acknowledgments:** The authors gratefully acknowledge the Australian Energy Market Operator (AEMO) for wholesale electricity market data, the Essential Services Commission (ESC) Victoria for Victorian Default Offer pricing data (2019-2025), and the Australian Bureau of Meteorology (BOM) for solar irradiance data. Solar Victoria provided detailed subsidy program documentation. The authors acknowledge Deakin University for computational facilities and MATLAB software access.

**During:** the preparation of this manuscript, the authors used Claude 4.6 Sonnet (Anthropic, 2025) for the purposes of manuscript editing, literature review assistance, and formatting of tables and figures. The authors have reviewed and edited the output and take full responsibility for the content of this publication. All numerical simulations, economic modelling, and scientific conclusions were independently performed by the authors.

**Conflicts of Interest:** Authors Ian Lilley and Jag Makam are affiliated with Zeco Australian Energy Solutions Pty Ltd., which supported and collaborated on this study. The company may potentially benefit from the outcomes presented in this paper. All efforts were made to ensure objectivity and scientific rigour.

## Abbreviations

VPP	Virtual Power Plant
V2G	Vehicle-to-Grid
G2V	Grid-to-Vehicle
EV	Electric Vehicle
BESS	Battery Energy Storage System
PV	Photovoltaic
NEM	National Electricity Market

NMI	Network Metering Installation
AEMO	Australian Energy Market Operator
D-FCAS	Demand-side Frequency Control Ancillary Service
BTM	Behind the Meter
MILP	Mixed Integer Linear Programming
ARMA	Autoregressive Moving Average
AIC	Akaike Information Criterion
MPC	Model Predictive Control
REMS	Residential Energy Management System
OCP	Open Charge Point Protocol
TOU	Time of Use
FiT	Feed-in Tariff
SoC	State of Charge
ADF	Augmented Dickey-Fuller
RMSE	Root Mean Square Error
MAE	Mean Absolute Error
AR	Autoregressive
MA	Moving Average
OLS	Ordinary Least Squares
RegD	Regulation D (signal)
AGC	Automatic Generation Control
KPI	Key Performance Indicator
RMS	Root Mean Square
DER	Distributed Energy Resources
LV	Low Voltage
AS/NZS	Australian Standard / New Zealand Standard
WDR	Wholesale Demand Response
AER	Australian Energy Regulator
ESC	Essential Services Commission
BOM	Bureau of Meteorology
AEST	Australian Eastern Standard Time

## References

1. Australian Energy Regulator *State of the Energy Market 2024*; AER: Canberra, ACT, Australia, 2024;
2. Electric Vehicle Council *State of Electric Vehicles: Australian EV Market Report 2024*; Electric Vehicle Council: Sydney, NSW, Australia, 2024;
3. Marzband, M.; Fouladfar, M.H.; Akorede, M.F.; others Framework for Smart Transactive Energy in Home-Microgrids Considering Coalition Formation and Demand Side Management. *Sustain. Cities Soc.* **2018**, *40*, 136–154.
4. Tushar, W.; Saad, W.; Poor, H.V.; Smith, D.B. Economics of Electric Vehicle Charging: A Game Theoretic Approach. *IEEE Trans. Smart Grid* **2012**, *3*, 1767–1778, doi:10.1109/TSG.2012.2211901.
5. Sortomme, E.; El-Sharkawi, M.A. Optimal Combined Bidding of Vehicle-to-Grid Ancillary Services. *IEEE Trans. Smart Grid* **2012**, *3*, 70–79, doi:10.1109/TSG.2011.2170099.
6. Liu, H.; Hu, Z.; Song, Y.; Lin, J. Decentralized Vehicle-to-Grid Control for Primary Frequency Regulation Considering Charging Demands. *IEEE Trans. Power Syst.* **2013**, *28*, 3480–3489, doi:10.1109/TPWRS.2013.2252029.
7. Hossain, M.A.; Pota, H.R.; Hossain, M.J.; Blaabjerg, F. Evolution of Microgrids with Converter-Interfaced Generations: Challenges and Opportunities. *Int. J. Electr. Power Energy Syst.* **2019**, *109*, 160–186, doi:10.1016/j.ijepes.2019.01.038.

8. Zhang, Q.; Dehghanpour, K.; Wang, Z.; Qiu, F.; Zhao, D. Multi-Agent Safe Policy Learning for Power Management of Networked Microgrids. *IEEE Trans. Smart Grid* **2021**, *12*, 1048–1062, doi:10.1109/TSG.2020.3034827.
9. Richardson, P.; Flynn, D.; Keane, A. Local Versus Centralized Charging Strategies for Electric Vehicles in Low Voltage Distribution Systems. *IEEE Trans. Smart Grid* **2012**, *3*, 1020–1028, doi:10.1109/TSG.2012.2185523.
10. Haupt, R.L. *Practical Genetic Algorithms*; 2nd ed.; John Wiley: Hoboken, N.J, 2004; ISBN 978-0-471-67175-6.
11. Box, G.E.P.; Jenkins, G.M.; Reinsel, G.C.; Ljung, G.M. *Time Series Analysis: Forecasting and Control*; Wiley series in probability and statistics; Fifth edition.; Wiley: Hoboken, New Jersey, 2016; ISBN 978-1-118-67492-5.
12. Mayer, M.J.; Gróf, G. Extensive Comparison of Physical Models for Photovoltaic Power Forecasting. *Appl. Energy* **2021**, *283*, 116239, doi:10.1016/j.apenergy.2020.116239.
13. Kempton, W.; Tomić, J. Vehicle-to-Grid Power Fundamentals: Calculating Capacity and Net Revenue. *J. Power Sources* **2005**, *144*, 268–279, doi:10.1016/j.jpowsour.2004.12.025.
14. Grau Unda, I.; Papadopoulos, P.; Skarvelis-Kazakos, S.; Cipcigan, L.M.; Jenkins, N.; Zabala, E. Management of Electric Vehicle Battery Charging in Distribution Networks with Multi-Agent Systems. *Electr. Power Syst. Res.* **2014**, *110*, 172–179, doi:10.1016/j.epsr.2014.01.014.
15. Nguyen, T.A.; Byrne, R.H.; Concepcion, R.J.; Gyuk, I. Maximizing Revenue from Electrical Energy Storage in MISO Energy & Frequency Regulation Markets. In Proceedings of the 2017 IEEE Power & Energy Society General Meeting; IEEE: Chicago, IL, USA, July 2017; pp. 1–5.
16. Sjeff Cobben; Bruno Gaiddon; Hermann Laukamp *Impact of Photovoltaic Generation on Power Quality in Urban Areas with High PV Population*; Intelligent Energy Europe: Brussels, Belgium, 2008;
17. Nambiar, J.; Yu, S.; Makam, J.; Trinh, H. Optimizing Power Sharing and Demand Reduction in Distributed Energy Resources for Apartments Through Tenant Incentivization. *Energies* **2025**, *18*, 4073, doi:10.3390/en18154073.
18. Wang, Y.; Mao, S.; Nelms, R.M. Online Algorithm for Optimal Real-Time Energy Distribution in the Smart Grid. *IEEE Trans. Emerg. Top. Comput.* **2013**, *1*, 10–21, doi:10.1109/TETC.2013.2273218.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.