Australia's National
Science Agency

Role of price-based aggregated control of CER in market dispatch

Nariman Mahdavi, Xinlin Wang, Vahid Aryai, Ghulam Mohy-ud-din, Subbu Sethuvenkatraman
May 2026

Citation

Mahdavi N, Wang X, Aryai V, Mohy-ud-din G, Sethuvenkatraman S (2026) Role of price-based aggregated control of CER in market dispatch. CSIRO, Australia.

Copyright

© Commonwealth Scientific and Industrial Research Organisation 2026. To the extent permitted by law, all rights are reserved and no part of this publication covered by copyright may be reproduced or copied in any form or by any means except with the written permission of CSIRO.

Important disclaimer

CSIRO advises that the information contained in this publication comprises general statements based on scientific research. The reader is advised and needs to be aware that such information may be incomplete or unable to be used in any specific situation. No reliance or actions must therefore be made on that information without seeking prior expert professional, scientific and technical advice. To the extent permitted by law, CSIRO (including its employees and consultants) excludes all liability to any person for any consequences, including but not limited to all losses, damages, costs, expenses and any other compensation, arising directly or indirectly from using this publication (in part or in whole) and any information or material contained in it.

CSIRO is committed to providing web accessible content wherever possible. If you are having difficulties with accessing this document please contact csiro.au/contact.

Acknowledgement of Country

CSIRO acknowledges the Traditional Owners of the lands, seas and waters of the are that we live and work on across Australia and pays its respects to Elders past and present. CSIRO recognises that Aboriginal and Torres Strait Islander peoples have made, and will continue to make, extraordinary contributions to Australian life including in cultural, economic, and scientific domains.

OFFICIAL

Contents

Acknowledgments vi

Executive summary vii

1 Introduction ix

Part I Market optimisation 1

2 NEMX.jl – A Physics-Based Co-Optimized Electricity and FCAS Market Modelling Framework 2

2.1 Overview of NEMX.jl 3

2.2 Synthetic NEM Dataset 5

2.3 Example Regional Pricing Outcomes 7

Part II Local Optimisation 9

3 HVAC Gym model 10

3.1 Simulation period 10

3.2 Extrapolation of HVAC Gym models 11

3.3 Providing diversity in building models 13

4 RL-Based Control Strategies for Smart Building Flexibility 15

4.1 The MDP Structure and the Critical Role of the Reward 15

4.2 Control Architectures and Algorithm Selection 16

5 Experimental Framework and Reward Mechanisms 18

5.1 Performance and Flexibility Analysis 21

5.2 Quantitative Assessment via the Flexibility Index (FI) 27

5.3 Price-demand mapping 29

Part III Whole system simulation 35

6 Impact of price-responsive resources on dispatched price 36

6.1 General network transformations 36

6.2 Connecting a flexible demand to sNEM2000 model 38

6.3 Python-Julia wrapper function 40

6.4 Simulation results 40

7 Conclusions 45

OFFICIAL

OFFICIAL

References 46

OFFICIAL

Figures

Figure 1 The proposed framework for studying the impact of price responsive resources on market dispatch vii

Figure 2 Conceptual architecture of NEMX.jl, illustrating the co-optimisation of the electricity and FCAS markets on a physics-based AC and HVDC network representation 4

Figure 3 Single-line diagram of the Synthetic NEM2000 dataset. Buses are plotted at their synthetic geographic coordinates, so the layout approximates the NEM footprint, and inter-node distance is approximately geographic; the angle between branches has no electrical meaning and branch length is not proportional to impedance. Branch colour denotes nominal voltage (kV) on a categorical scale with five entries: 132, 220, 275, 330, 500 kV as a continuous ramp. Green segments are two-winding transformers: a transformer is modelled as a branch between the HV and LV buses of a substation, and because those buses share coordinates it appears as a short stub rather than a point. Red segments are the three HVDC interconnectors: Directlink (NSW–QLD), Murraylink (SA–VIC) and Basslink (VIC–TAS). 7

Figure 4 Energy prices at regional reference nodes 8

Figure 5 FCAS prices at regional reference nodes 8

Figure 6 A) HVAC Gym model components, B) Overview of the workflow for RL-based HVAC control using HVAC Gym. 10

Figure 7 The warmest week of 2025 in Newcastle on median basis 11

Figure 8 Training Newcastle Gym models on odd days and test on even days. 12

Figure 9 True and predicted chiller electric power for Newcastle site trained on data from 8/2023-3/2024 and tested on 11/2024-3/2025. 12

Figure 10 Training Newcastle Gym models using 95th percentile of data and test them on the remaining extreme data. 13

Figure 11 Train on Newcastle, Test on Clayton using the second approach 14

Figure 12 Train on Newcastle, Test on Clayton using the first approach 14

Figure 13 Representative 10-day external signals used in the RL evaluation: (a) raw AEMO NSW1 spot electricity price and (b) outdoor air temperature, both shown at a 10-minute control interval. 21

Figure 14 Performance of General RL under price-weighted reward: Comparison of (a) Wide Thermal Boundaries (T1) and (b) Narrow Thermal Boundaries (T2). 22

Figure 15 Performance of Constrained RL under price-weighted reward: Comparison of (a) Wide Thermal Boundaries (T1) and (b) Narrow Thermal Boundaries (T2). 23

Figure 16 Performance of General RL under outdoor-weighted reward: Comparison of (a) Wide Thermal Boundaries (T1) and (b) Narrow Thermal Boundaries (T2). 25

Figure 17 Performance of Constrained RL under cost-only reward: Comparison of (a) Wide Thermal Boundaries (T1) and (b) Narrow Thermal Boundaries (T2). 26

Figure 18 Aggregated Flexibility Map: Average performance across all 8 tested control strategies and thermal boundary conditions 27

Figure 19 Performance of the developed sequence to sequence and one to one models for predicting total load and cooling usage. The top plots show the one-to-one models while the bottom plots show results for sequence to sequence models. 33

Figure 20 Scaling factors for all the loads in NEMX.jl and the resulting calculated dispatch price, when we run NEMX in timeseries mode. 36

Figure 21 Price-demand bids of scheduled loads in NEM, taken from AEMO's guide to scheduled loads, July 2021 37

Figure 22 Flexible chiller demand and total demand of Newcastle Energy Centre in a typical week in February 2025 39

Figure 23 NSW section of the sNEM 2000 model and node 690 as Newcastle site representation 39

Figure 24 Impact of ToU-triggered non-scheduled flexible loads on dispatch price 41

Figure 25 Impact of being schedulable load at node 690 on dispatch price with total demand multiplying factor 10 42

Figure 26 Impact of being schedulable load at node 690 on dispatch price with total demand multiplying factor 150 42

Figure 27 Impact of being schedulable load at node 605 on dispatch price with total demand multiplying factor 150 43

Figure 28 Same scenario as in Figure 25, but with different generation cost bidding curves based on fuel types 44

Tables

Table 1 10 FCAS markets modelled in NEMX 4

Table 2 Spot price regions in NEMX 5

Table 3 Structural characteristics of the S-NEM2000 dataset 6

Table 4 Extrapolation performance of HVAC Gym models for predicting chiller electrical power. 11

Table 5 Summary of Control Strategies and Experimental Scenarios 20

Table 6 Permutation-test results for the significance and detectability of the modelled price response 30

Table 7 Results for 1-1 regression between price, outside air temperature and hour of day and site cooling usage 31

Table 8 Results for 1-1 regression between price, outside air temperature and hour of day and site total load. 31

Table 9 Results for sequence to sequence (24 of each input to 24 of target) regression between price, outside air temperature and hour of day and site cooling usage. 32

Table 10 Results for sequence to sequence (24 of each input to 24 of target) regression between price, outside air temperature and hour of day and site total load..... 32

Table 11 Slopes and bias terms for the one to one cooling usage models. 33

Table 12 Slopes and bias terms for the one to one total load models. 34

Table 13 Slopes and bias terms for the sequence to sequence cooling usage models. 34

Table 14 Slopes and bias terms for the sequence-to-sequence total load models. 34

Acknowledgments

Australian Research for power system transformation (AR-PST)

Executive summary

In this project, we look at the whole electricity system from price responsive CERs located at the low-voltage (LV) level to their aggregations at the medium-voltage (MV) level, and a wholesale energy and FCAS market determining the wholesale spot price based on a synthetic National Electricity Market (NEM) model. In our proposed framework in Figure 1, we investigate the impact of spot price on the behaviour of an aggregation of price responsive resources (PRRs) through the identification of simple Flexibility Functions (FFs). FF is a map between spot price and aggregated demand of flexible loads in response to the price signal.

While the proposed framework in this project is applicable to any type of flexible loads, our focus is on the flexibility from HVAC systems in commercial buildings. The goal of a local optimisation in a building at the LV level is to minimise the chiller electricity demand or cost while satisfying thermal comfort and local grid requirements. Since the electricity demand of a single building is not at the required level to be connected to the synthetic NEM model, we introduce a total demand multiplying factor that magnifies the demand profile. The FF is then found based on an ensemble of such magnified buildings each of which using a different set of local optimisation goal and/or constraint. We will finally investigate how a change in dispatchable aggregation of flexibility can alter the spot price.

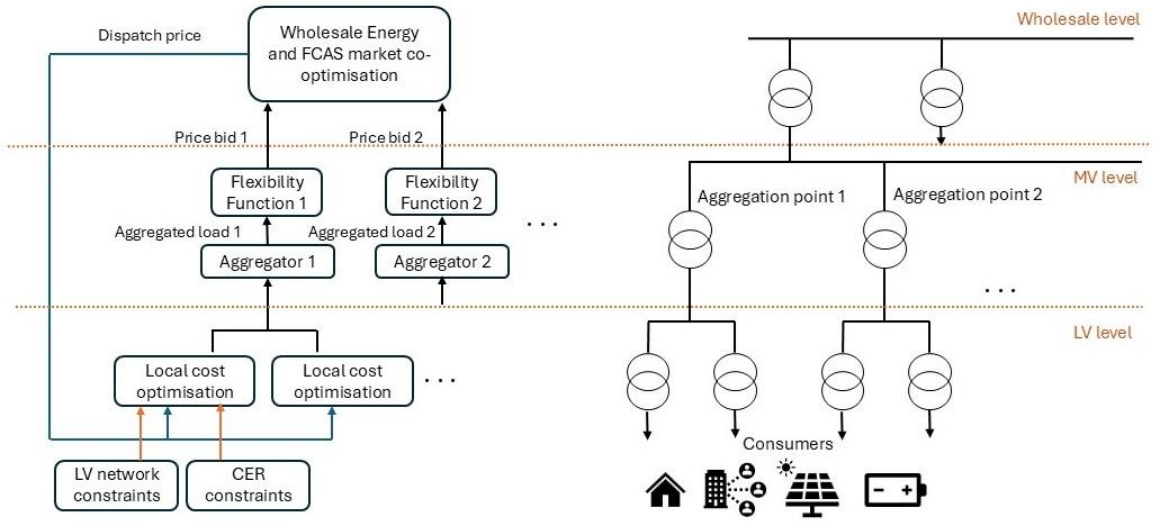


Figure 1 The proposed framework for studying the impact of price responsive resources on market dispatch

In summary, this research includes

- Implementation of the overall framework in Figure 1, where the Julia package NEMX.jl for energy and FCAS market co-optimisation is combined with Python Gym package for local RL-based chiller optimisation. The integration of these two environments includes running some standard testing scenario to make sure consistent results.
- Improving the existing NEMX package for timeseries analysis. While the original NEMX package runs snapshot optimisation scenarios, it cannot be run in timeseries mode required for flexibility analysis.

- Understanding the impact of variability and uncertainty of flexible loads on the dispatch price. Here, we investigated the impact of time-varying load profile and synchronous change in demand on the spot price.
- Investigating a potential spot price reduction by scheduling an aggregation of flexible loads and understanding the impact of their sizes, capacity-price bidding values, and locations.
- Understanding the impact of synchronous local optimisation (in an ensemble) on the dispatch price. This is simulated by training RL agents to receive AEMO's historical spot price signal and minimise the cost of energy at a building, while minimising the thermal discomfort and grid constraints.
    a) To diversify the experiments, a range of optimisation goals and constraints have been considered.
    b) Only a fixed operating envelop (OE) has been considered. We did not consider dynamic operating envelope.

1 Introduction

In an electricity market, a system operator schedules energy production and consumption, and deploy spinning or non-spinning reserve, e.g., to accommodate intermittent renewable energy sources. Australian electricity system is changing with high levels of Consumer Energy Resources (CER) ranging from batteries, PV and EV, where AEMO's ISP 2024 predicts coordinated CER storage will be nearly 50% of dispatchable capacity in the NEM by 2050. CERs, however, comes with uncertainty from environmental factors that may impact generation/consumption patterns to uncertainties on how consumers may respond to change in price, e.g., Time of Use tariffs. To mitigate the risks and utilise the opportunities, Australian government has developed CER roadmap and AEMC is currently carrying out rule changes for Integrating Price Responsive Resources (IPRR) enabling participation of CER in the National Electricity Market (NEM).

While PRR (price response resources) are capable of shaping demand in response to market signals, how their implicit flexible response will impact capacity allocation is not known. For instance, CER operation within the optimisation of the wholesale market and options for CER dispatching have identified by AEMO as highest priority focus areas. This includes incorporating Dynamic Operating Envelopes (DOEs) in the optimisation. Existing economic dispatch algorithms proposed in literature have mostly focused on the improvement of algorithm performance and consideration of more operational constraints such as spinning reserve and generator ramp rates. They did not dynamically engage an aggregation of CERs with time-varying price mechanisms such that the end-users are motivated to control their CERs to reduce electricity bill. Additionally, if these resources are effectively integrated into the wholesale market dispatch process, they can deliver essential power system services, and hence, i) reduce dispatch costs in the NEM, ii) reduce energy prices, and iii) support with reliability of grid.

To make that happen and with the increasing penetration of CERs, a coordination between Distribution Network Service Providers (DNSPs) and AEMO as the market operator is required. While the need for such coordination is imminent, there is not yet an interaction mechanism between AEMO and DNSPs to facilitate that. In this regard, there are two major considerations: i) CER need to be able to deliver these services while maintaining network reliability within network constraints, ii) Consumers who are the owner of these assets need to have simple ways to allow their resources to deliver grid services and be financially compensated in the easiest form.

Use of price signal to coordinate CER is considered an attractive approach due to its simplicity of implementation but has not been studied well. Although activities like Project Edith have attempted to evaluate the potential of price responsive CER on network services, price responsive CER coordination and its effect in delivering market services is not understood. This project aims to address this gap by evaluating ways that coordinated CER can participate in market dispatch while honouring network constraints and delivering energy bill savings to consumers.

This study will investigate the impact of flexible loads variability and uncertainty, CERs local site-level cost optimisation, and DOE constraints on the optimal scheduled/dispatch price provided by our NEM dispatch tool that co-optimises wholesale Energy and FCAS markets for a given FCAS product. In this study, we focus on the flexibility of HVAC loads in a commercial building with a

local optimiser minimising the cost of electricity at the LV level given the thermal comfort and local network constraints provided by a DOE signal. We then consider an ensemble of such commercial buildings to reach to the aggregator level sitting at the MV scale. The flexibility function, which is a map between price signal and demand will then be calculated at the aggregator level. To ensure scalability, flexible loads will be modelled by data-driven machine learning (ML) tools and local cost optimisation will be based on Reinforcement Learning (RL) optimisation.

Note while our focus is on HVAC flexibility, our proposed framework is still valid for any other type of flexible CER if their control is happening at the site level. Direct control of the CERs by an aggregator is out of the scope of this project. Saying that and although we did not consider any load variation by an aggregator, the implicit flexible load variation through the alteration of the price signal can be considered by the aggregator. This future work will require a second optimiser at the MV level to adjust the price signal, using the price-demand flexibility function, such that the aggregated load profile is following a desired load profile.

Part I Market optimisation

2 NEMX.jl – A Physics-Based Co-Optimized Electricity and FCAS Market Modelling Framework

The Australian Energy Market Operator-operated National Electricity Market (NEM) is one of the largest interconnected electricity markets globally, spanning long transmission distances and integrating diverse generation technologies across multiple regions. Electricity dispatch in the NEM is centrally coordinated through a real-time market clearing process, where generation offers are optimized to meet demand at minimum cost while respecting system security constraints. A defining feature of the NEM is the co-optimization of energy and Frequency Control Ancillary Services (FCAS) within the dispatch engine (NEMDE). This co-optimization ensures that both energy production and frequency regulation requirements are simultaneously satisfied in each dispatch interval, enabling efficient allocation of generation resources across competing objectives. The FCAS markets in the NEM are essential for maintaining system frequency stability in response to disturbances such as generation outages or sudden load changes. These services operate across multiple timeframes from fast contingency response (seconds) to slower regulation services and are co-optimized with energy dispatch using shared capacity constraints and system-wide requirements.

The rapid transformation of the NEM towards a high-renewable, inverter-based power system introduces significant challenges for both market operation and system security. Increased penetration of variable renewable energy sources, such as wind and solar, reduces system inertia and increases reliance on fast frequency response services. As a result, the role of FCAS markets has become increasingly critical in ensuring secure operation under uncertainty. At the same time, the expansion of High Voltage Direct Current (HVDC) interconnectors is reshaping power flows across regions. HVDC links, such as Basslink and other interconnectors, provide controllable power transfer capabilities but introduce additional complexity in market clearing due to non-linear losses, controllability, and interaction with AC network constraints. Traditional market clearing models often approximate or neglect these characteristics, leading to inefficiencies and suboptimal dispatch outcomes. Furthermore, the increasing need for real-time, security-constrained and network-aware market clearing has highlighted limitations in simplified DC-based or zonal models. Accurate representation of AC power flow physics, voltage constraints, and converter behaviour is essential to ensure feasibility and reliability in operational decisions, particularly under stressed system conditions.

To address these challenges, NEMX.jl[1] has been developed as an open-source Julia-based package for co-optimized electricity and frequency control ancillary services (FCAS) market clearing. Built on the Julia-JuMP framework, the package enables physics-based modelling of both AC and HVDC

[1] Available [Online]: https://github.com/csiro-internal/ensys-NEMX.jlhttps://data.csiro.au/collection/csiro%3A66378v2, https://doi.org/10.25919/v33w-8011

networks. NEMX.jl integrates detailed market mechanisms with network constraints, providing a comprehensive platform for studying operational and planning problems in the NEM.

NEMX.jl builds on an established line of open network models of the NEM. The underlying Synthetic-NEM-2000bus network representation and the AC and AC–DC optimal power flow formulations on which the market clearing engine operates were developed previously and are documented in [3]–[5]. The capabilities developed for the present study are: (i) the co-optimised energy and FCAS market clearing formulation, covering the ten FCAS markets of Table 1 and their trapezium enablement constraints for a time-series simulation; (ii) addition of time-series data for solar, wind, and load demand; (iii) addition of the Newcastle site flexible and fixed load profiles as new participants acting as a dispatchable load; (iv) recovery of regional reference node (RRN) energy prices, regional FCAS prices and nodal locational marginal prices from the co-optimised solution for the time-series simulation; and (v) the Python–Julia interface described in Section 6.3, which allows an externally simulated, price-responsive demand to be represented inside the market clearing loop. These extensions are available in NEMX.jl [6].

2.1 Overview of NEMX.jl

NEMX.jl as shown in Figure 2 is designed to replicate key aspects of the NEM dispatch engine, incorporating both market-based and physics-based formulations. It utilizes the Synthetic-NEM-2000bus dataset[2], enabling large-scale, realistic simulations of the Australian power system. The key capabilities of this tool are listed as follows:

- Power and FCAS Bids in terms of price and quantity bands.
- FCAS trapezium constraints including enablement limits, bounds, and availability.
- Dispatchable loads participation in power and FCAS markets.
- Dispatchable bidirectional units such as storage systems and aggregators.
- Regional reference node spot pricing for power and FCAS.
- Power and FCAS unit power dispatch.
- Nodal locational marginal pricing (LMP).
- Marginal Loss Factor (MLF) modelling.
- Regional FCAS targets.
- Transmission system and AC as well as HVDC interconnector flows.

[2] R. Heidari, M. Amos and F. Geth, "An Open Optimal Power Flow Model for the Australian National Electricity Market," 2023 IEEE PES Innovative Smart Grid Technologies - Asia (ISGT Asia), Auckland, New Zealand, 2023, pp. 1-5, doi: 10.1109/ISGTAsia54891.2023.10372618.
Available [Online]: https://github.com/csiro-energy-systems/Synthetic-NEM-2000bus-Data

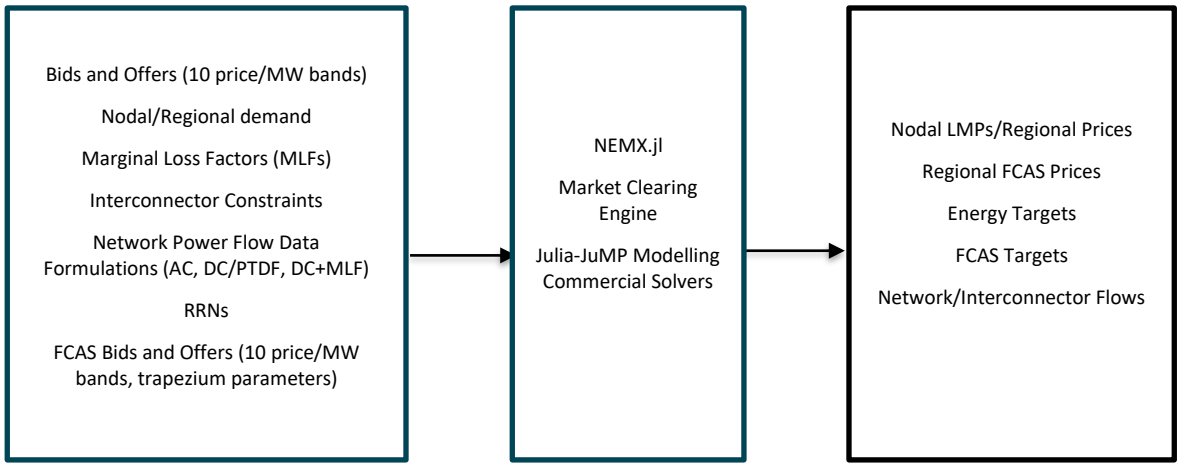


Figure 2 Conceptual architecture of NEMX.jl, illustrating the co-optimisation of the electricity and FCAS markets on a physics-based AC and HVDC network representation

The FCAS markets and spot price regions modelled in the NEMX.jl are listed in Table 1 and Table 2, respectively.

Table 1 10 FCAS markets modelled in NEMX

Service Name	Name Tag	Response Time Description
Regulation Raise	RReg	Minor frequency drop – primary frequency response
Regulation Lower	LReg	Minor frequency rise – primary frequency response
Very Fast Raise	R1S	Major frequency drop – 1s
Very Fast Lower	L1S	Major frequency rise – 1s
Fast Raise	R6S	Major frequency drop – 6s
Fast Lower	L6S	Major frequency rise – 6s
Slow Raise	R60S	Stabilise after major drop – 60s
Slow Lower	L60S	Stabilise after major rise – 60s
Delayed Raise	R5M	Return to normal operating band – 5m
Delayed Lower	L5M	Return to normal operating band – 5m

Table 2 Spot price regions in NEMX

Region Code	Region Name
area 1	New South Wales (NSW)
area 4	Queensland (QLD)
area 2	Victoria (VIC)
area 3	South Australia (SA)
area 5	Tasmania (TAS)

2.2 Synthetic NEM Dataset

The increasing need for high-fidelity modelling of large-scale power systems has driven the development of synthetic datasets that replicate the structural and operational characteristics of real-world networks while preserving confidentiality. Some examples from US and Europe are available in references [7], [8], and [9]. In the context of the Australian National Electricity Market (NEM), the Synthetic NEM dataset provides a comprehensive, publicly usable representation of the transmission system, enabling advanced research in system planning, operation, and market analysis.

Initially, a synthetic phasor-domain and EMT real-time digital simulation (Hypersim) dataset was developed based on a PSS®E network model [1], [2]. The dataset was subsequently extended for power network optimisation studies [3], with HVDC interconnectors later incorporated in [4]. The Synthetic NEM dataset is designed to emulate the topology, electrical characteristics, and geographic distribution of the real NEM, which spans multiple states and extensive transmission corridors. It comprises approximately two thousand buses and 3,800 transmission circuits, together with several hundred generators and aggregated loads, capturing the scale and operational complexity of the interconnected Australian power system.

The dataset is synthetic rather than anonymised. Bus identifiers, substation names and coordinates do not correspond to real assets; branch impedances, ratings and voltage assignments are representative of the relevant asset class rather than as-built values; and demand is aggregated to regional sub-area centroids. What is retained is the network topology and degree of meshing, the voltage-level hierarchy, the regional boundaries and reference nodes, the HVDC interconnector arrangement, and the aggregate regional demand and installed generation capacity. This combination is what allows the dataset to support network-constrained market studies without disclosing confidential infrastructure detail [1]–[3]. Commonly used variants, such as the S-NEM2000 dataset shown in Figure 3, include approximately 2,000 AC buses together with detailed representations of transmission infrastructure across all major NEM regions. Figure 3 shows the resulting network, in which the five NEM regions, the radial coastal corridors and the three HVDC links are all identifiable. The structural characteristics of S-NEM2000 are summarized in Table 3.

A key advantage of the Synthetic NEM dataset lies in its ability to support network-constrained and physics-based analyses. Unlike simplified test systems, it incorporates realistic impedance

parameters, voltage levels, and inter-regional connections, enabling studies that require accurate representation of power flows, losses, and operational limits. This makes it particularly suitable for applications such as optimal power flow (OPF), security-constrained dispatch, and transmission expansion planning [5]. In addition to AC network representation, extended versions of the dataset include HVDC interconnectors and converter models [4], reflecting the growing importance of power electronics in the NEM. These extensions enable the study of hybrid AC–DC system behaviour, including controllable power flows, converter losses, and interactions between synchronous and inverter-based resources. Such capabilities are essential for analysing modern grid challenges associated with renewable integration and inter-regional energy exchange.

The dataset is also widely used in conjunction with time-series inputs, such as demand profiles, renewable generation traces, and market data, allowing for time-sequential and stochastic analyses. This facilitates the evaluation of system performance under varying operating conditions, including peak demand events, renewable variability, and contingency scenarios. As a result, the Synthetic NEM dataset provides a robust foundation for both operational studies and long-term planning assessments. Importantly, the Synthetic NEM dataset maintains a balance between data accessibility and realism. By abstracting sensitive infrastructure details while preserving key system characteristics, it enables collaboration between academia, industry, and policymakers. This has supported the development of open-source tools, benchmarking studies, and reproducible research in areas such as market design, system security, and resilience under extreme events [6], [5].

In summary, the Synthetic NEM dataset serves as a critical enabler for high-resolution modelling and analysis of the Australian power system. Its scalability, flexibility, and compatibility with modern optimization frameworks make it an essential resource for advancing research in power system engineering, particularly in the context of transitioning towards low-carbon and highly interconnected electricity networks.

Table 3 Structural characteristics of the S-NEM2000 dataset

Component	Quantity	Notes
AC buses	≈ 2,000	Synthetic buses, identifiers, substation names and coordinates do not map to real NEM assets; topology, voltage hierarchy and regional structure are preserved [1]–[3]
AC transmission lines and transformers	≈ 3800	Includes 132 kV, 220 kV, 275 kV, 330 kV, and 500 kV circuits with two- and three-winding transformers
HVDC interconnectors	3	Basslink (VIC to TAS), Murraylink (SA to VIC), Directlink (NSW to QLD)
Generators	≈ 261	Baseline 2024 set
Loads	≈ 912	Distributed across regional sub-area centroids
Regions	5	NSW, QLD, VIC, SA, TAS

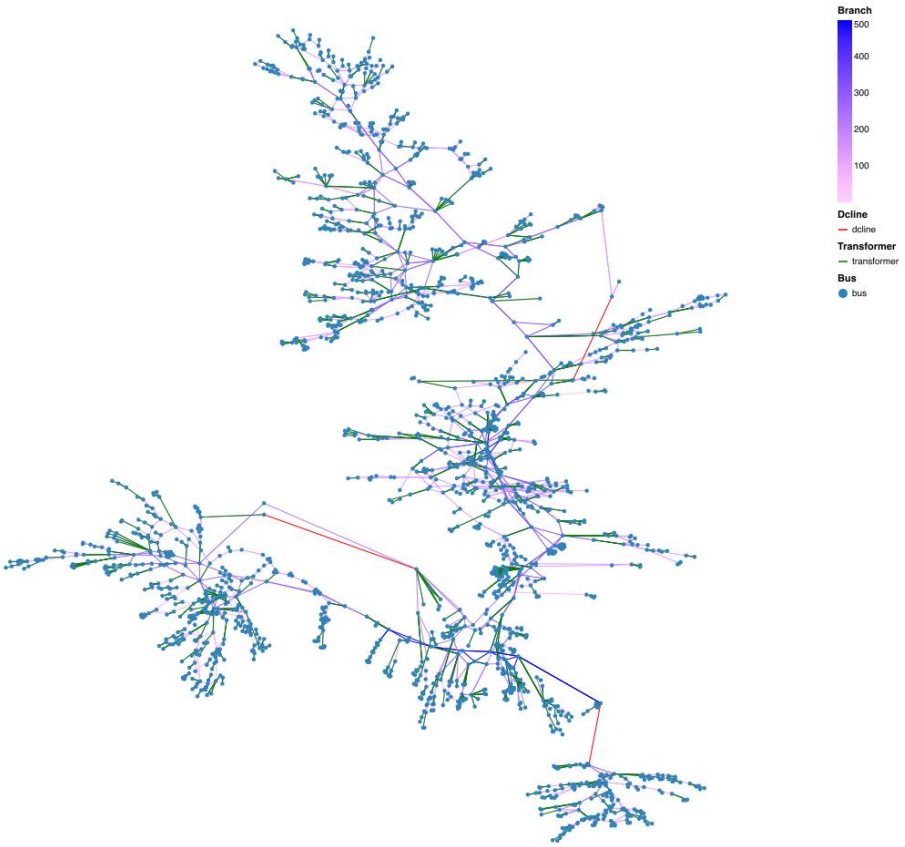


Figure 3 Single-line diagram of the Synthetic NEM2000 dataset. Buses are plotted at their synthetic geographic coordinates, so the layout approximates the NEM footprint, and inter-node distance is approximately geographic; the angle between branches has no electrical meaning and branch length is not proportional to impedance. Branch colour denotes nominal voltage (kV) on a categorical scale with five entries: 132, 220, 275, 330, 500 kV as a continuous ramp. Green segments are two-winding transformers: a transformer is modelled as a branch between the HV and LV buses of a substation, and because those buses share coordinates it appears as a short stub rather than a point. Red segments are the three HVDC interconnectors: Directlink (NSW–QLD), Murraylink (SA–VIC) and Basslink (VIC–TAS).

2.3 Example Regional Pricing Outcomes

This section presents representative regional pricing outcomes obtained from the co-optimized market clearing model for scenario 3 available in script ex_opfcas_arpst.jl in NEMX.jl. The results presented in Figures 4 and 5 illustrate spatial variations in electricity prices across the NEM regions, reflecting underlying network constraints, generation dispatch patterns, and inter-regional power flows. Differences in regional prices are driven by congestion, marginal loss factors, and the availability of low-cost generation, highlighting the importance of network-aware market modelling. The separation of the NSW price from the southern regions reflects a binding import limit into NSW under this scenario; the southern regions clear against low-cost marginal generation and price together at $22.3/MWh.

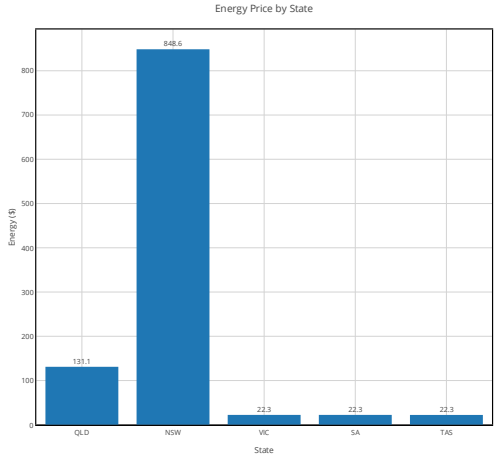


Figure 4 Energy prices at regional reference nodes

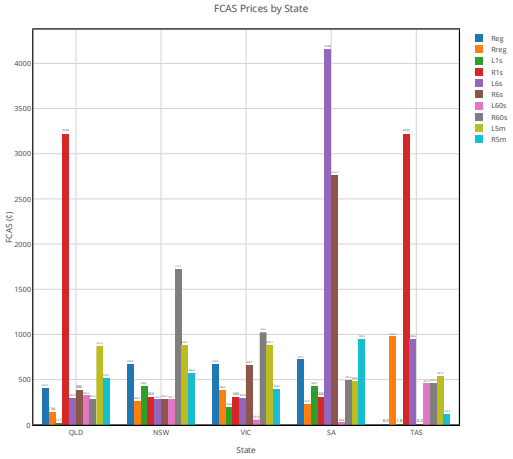


Figure 5 FCAS prices at regional reference nodes

This chapter has presented the NEMX.jl framework for physics-based co-optimised electricity and FCAS market modelling, documented its underlying synthetic-NEM dataset, and demonstrated its application to the analysis of regional energy and FCAS prices in market dispatch. The NEMX.jl framework provides the open-source, network-aware analytical foundation required to support the design and assessment of post-2027 market reforms in the NEM and is positioned as the operational counterpart to the NEMDE in this project to study the price-responsive CER participation in the NEM as detailed in the next chapters.

Note that NEMX.jl does not reproduce AEMO's constraint library of network and system-security constraint equations, does not model unit commitment or the full NEMDE tie-breaking and pricing run structure, and operates on a synthetic rather than the actual network.

Part II Local Optimisation

3 HVAC Gym model

Since in this project we are focusing on the price-based local flexibility of HVAC in a commercial building, we use CSIRO's HVAC Gym model. HVAC Gym represents the coupled dynamics of both the building thermal environment and the HVAC system, and not the HVAC system alone. It provides a data-driven modelling framework for capturing HVAC-building interactions at scale. The framework wraps a set of coupled machine-learning models each modelling a part of the HVAC-building interaction dynamics using Open AI's Gym. This way HVAC Gym models provide a standard framework for implementing data-driven control approaches, such as Reinforcement Learning (RL) agents, to control buildings' HVAC electricity demand. A full description of the methodology behind CSIRO's HVAC Gym can be found in the works of [13] with its successful real-world implementation [11]. An overview of the HVAC Gym and the workflow for RL-based HVAC control using the Gym given in Figure 6. In this project, we use the HVAC gym models for CSIRO Newcastle Energy Centre to implement the RL agents for controlling buildings HVAC as discussed in Section 4.

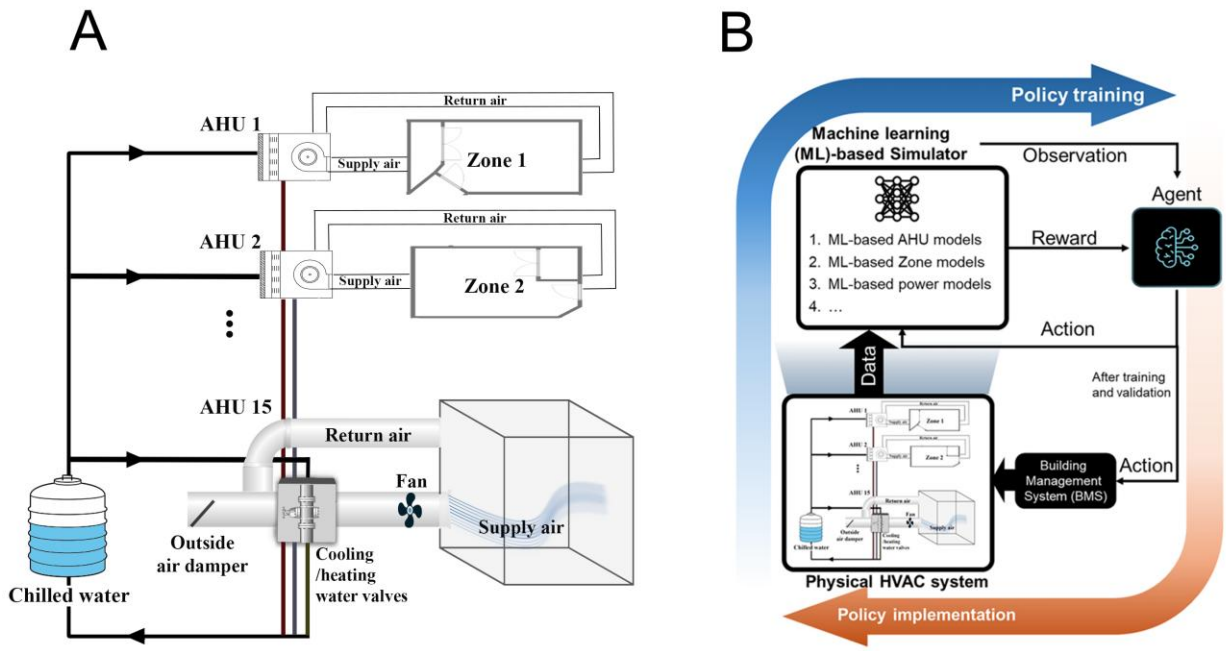


Figure 6 A) HVAC Gym model components, B) Overview of the workflow for RL-based HVAC control using HVAC Gym.

3.1 Simulation period

In this research, all the simulations are focused on a week of 2025 where the median temperature were the highest annually in Newcastle with the aim of considering a constrained gird period. Figure 7 shows the outside air temperature throughout that week and shows chiller power and total building load of the Newcastle site.

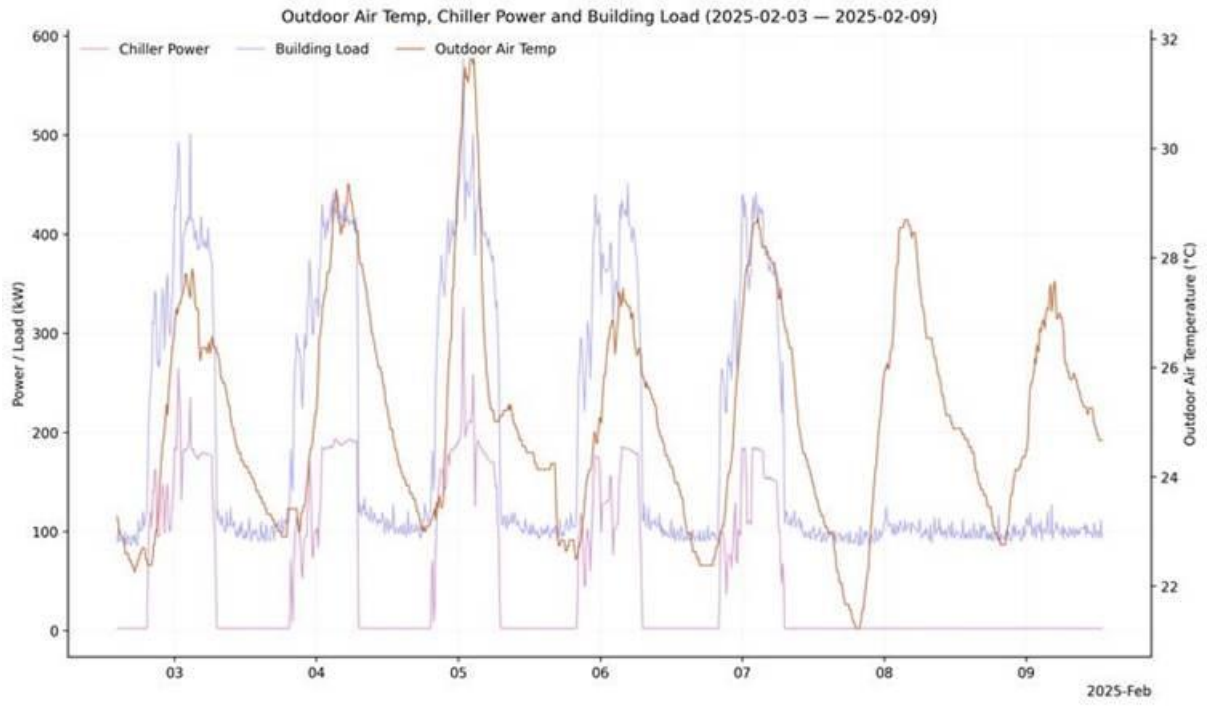


Figure 7 The warmest week of 2025 in Newcastle on median basis

3.2 Extrapolation of HVAC Gym models

Since simulation of the warmest week of the 2025 was intended in this report and given that the original HVAC gym models were developed and tested for 2023-2024, we examined the extrapolation capabilities of the models to make sure they are fit for such an outlier week simulation. To achieve this aim, we tried the following three evaluation approaches with their results listed in Table 4. As it can be seen both Newcastle and Clayton buildings Gym models provide more than 0.7 R-squared, in average, and hence, suitable for simulation.

1 Training Gym models on odd days and test on even days (see Figure 8), and

2 Training Gym models using data from a year and testing them using data from a different time/year (see Figure 9).

3 Training Gym models using 95th percentile of data and test them on the remaining extreme data. In other words, it checks if a model trained on typical target values (e.g., chiller power) can generalise to unusually low or high target values that were excluded from training (see Figure 10).

Table 4 Extrapolation performance of HVAC Gym models for predicting chiller electrical power.

Approach	Building	Train period	Test period	R-Squared	RMSE
1	Newcastle	11/24-3/25 (odd days)	11/24-3/25 (even days)	0.71	0.04
1	Clayton	11/22-3/23 (odd days)	11/22-3/23 (even days)	0.87	1.58

2	Newcastle	8/24-3/25	11/23-3/24	0.73	0.048
2	Clayton	8/22-3/23	11/21-12/21	0.5	4.305
3	Newcastle	11/24-3/25 (95th percentile)	11/24-3/25 (the rest)	0.73	0.076
3	Clayton	11/22-3/23 (95th percentile)	11/22-3/23 (the rest)	0.83	5.17

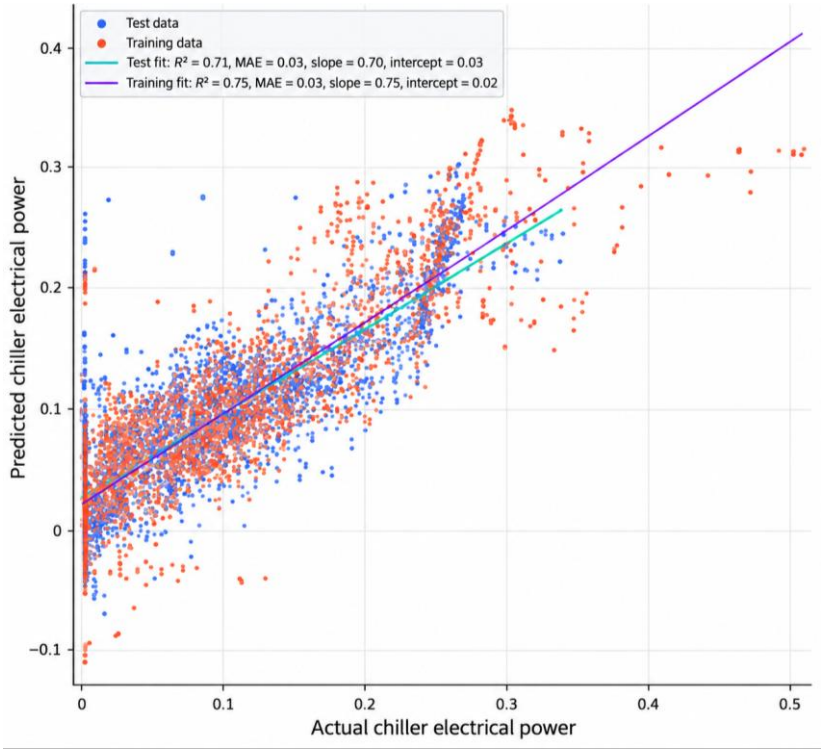


Figure 8 Training Newcastle Gym models on odd days and test on even days.

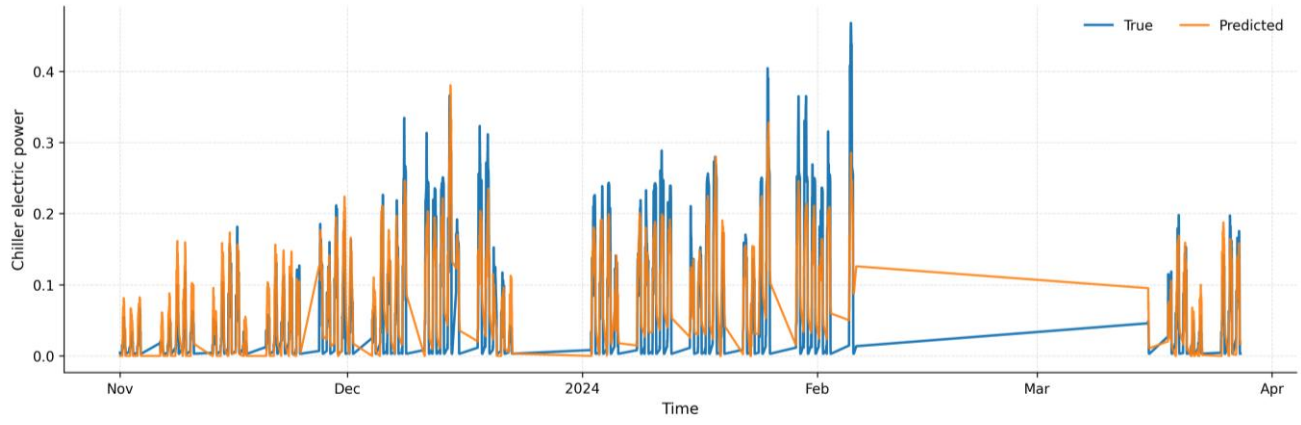


Figure 9 True and predicted chiller electric power for Newcastle site trained on data from 8/2023-3/2024 and tested on 11/2024-3/2025.

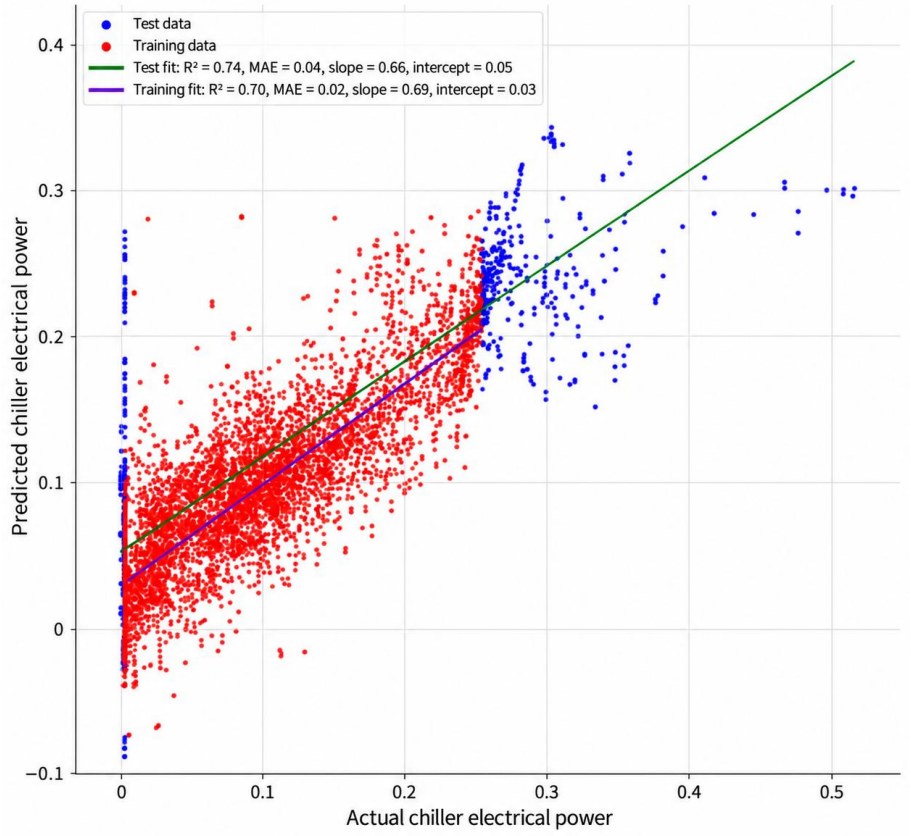


Figure 10 Training Newcastle Gym models using 95th percentile of data and test them on the remaining extreme data.

3.3 Providing diversity in building models

To investigate the potential for improving model diversity when only two building models were available, an exploratory building-swapping strategy was examined. In this strategy, the Gym model developed for one site was tested using data from the other site, for example, applying the Newcastle Gym model to Clayton data, and vice versa. This approach was intended to assess if cross-site model use could expand the range of building responses available for training or evaluation. One problem to implement this building swapping, is that there are small differences in data availability between Clayton and Newcastle site. For example, Fan speed data that used in developing Newcastle Gym might not be available for Clayton site. Two different approaches used to address that data inconsistency:

1) Developing Gym models for each site only using a mutually available data between sites and ignore the data that exist for one site and not for the other site, and

2) Developing Gym models based on the original implementation but considering a constant value for the parameters those do not exist in a site. For example, assuming a constant 25% fan speed for Clayton site which does not have this data.

The figures below demonstrate the effectiveness of each of these two approaches. Figure 11 shows the underlying chiller power model of the Newcastle Gym performance if developed using the second approach and tested using Clayton site's data. Figure 12 demonstrates the results of the first approach described above. It can be concluded from metrics (for example a 0.7 vs 0.3 R-

squared) that the second approach is better than the first one. This suggests that building swapping may be a useful direction for future work. However, this approach was not used in the final study. Therefore, the analysis is reported only as an indication of potential and not as part of the final research methodology.

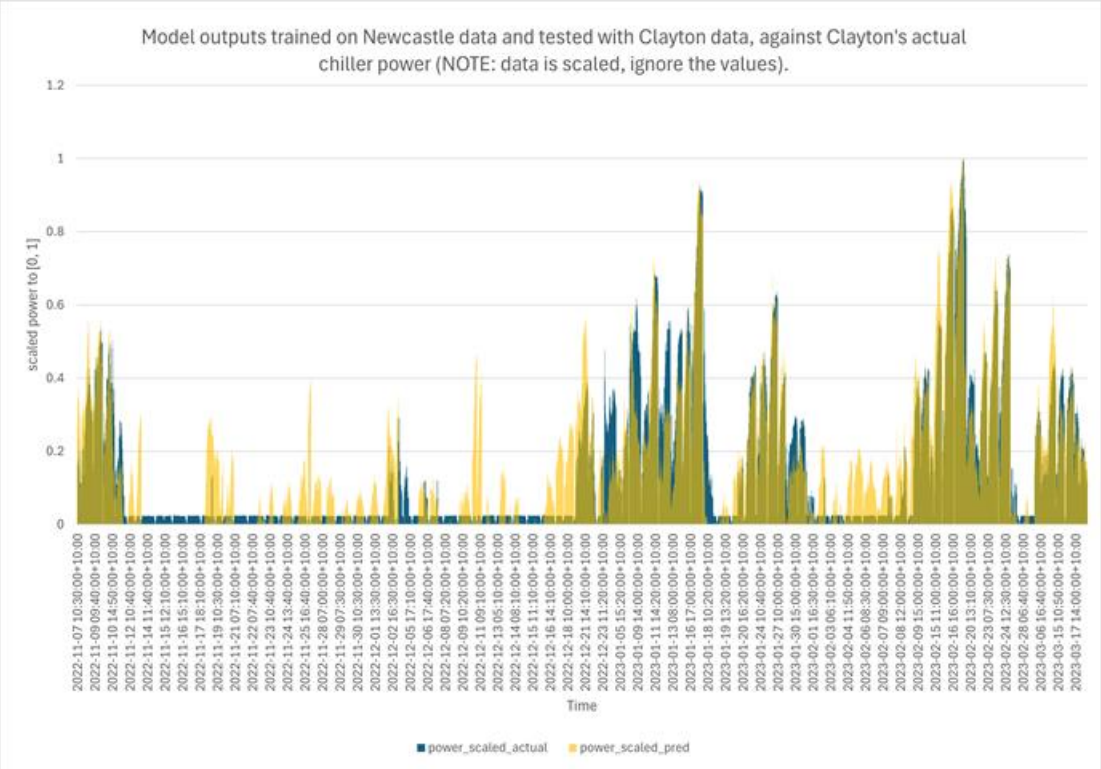


Figure 11 Train on Newcastle, Test on Clayton using the second approach

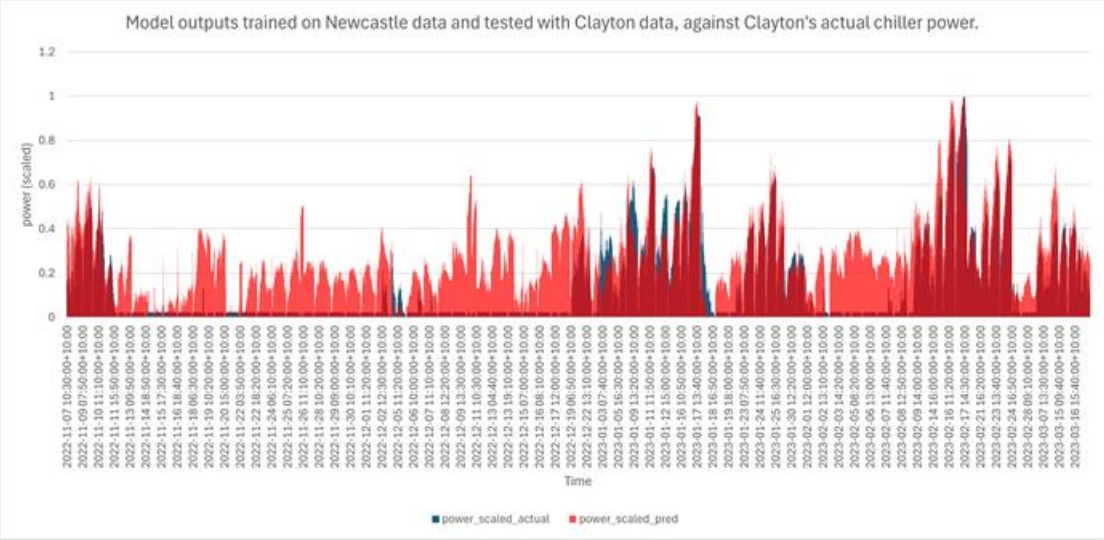


Figure 12 Train on Newcastle, Test on Clayton using the first approach

4 RL-Based Control Strategies for Smart Building Flexibility

The transition toward smart grid integration and sustainable building operations requires control systems capable of managing highly stochastic environments. Traditional heuristic and Rule-Based Control (RBC) strategies, while robust, are inherently static. They struggle to optimize the non-linear thermodynamics of building environments under time-varying disturbances, such as dynamic electricity pricing, fluctuating weather conditions, and unpredictable occupancy loads.

To address these limitations, Reinforcement Learning (RL) has emerged as a promising data-driven alternative. By formulating HVAC control as a Markov Decision Process (MDP), RL can learn adaptive control policies directly through interaction with the environment. Unlike conventional RBC, RL can capture complex system dynamics, respond to uncertainty, optimize long-term performance, and balance competing objectives such as energy consumption, operational cost, and occupant comfort without requiring explicit system models or handcrafted control rules.

4.1 The MDP Structure and the Critical Role of the Reward

To understand how different RL strategies manage building operations, it is necessary to pinpoint the underlying structure of the MDP, which is defined by three core interactions:

- **State Space (Observations):** The dynamic variables the agent observes at any given time step, such as indoor temperature, outdoor weather, and real-time electricity prices.
- **Action Space:** The control decisions available to the agent, such as adjusting temperature setpoints or modulating HVAC power.
- **Reward Signal:** The numerical feedback returned by the environment after an action is executed.

The reward mechanism is the fundamental driver of the RL framework. An agent possesses no inherent knowledge of "comfort" or "cost"; its sole mathematical objective is to discover a policy that maximizes cumulative rewards over time. Consequently, the formulation of the reward function dictates the entire behavioural outcome of the system. If the reward mechanism is poorly designed, the agent will exploit loopholes, leading to suboptimal or unsafe operations. Conversely, an effectively structured reward allows the system to dynamically balance complex, competing objectives.

By leveraging the dynamic nature of RL, it offers distinct operational advantages:

- **Model-Free Flexibility:** RL agents can learn optimal control policies directly from data and environmental interaction, bypassing the need to develop complex, computationally expensive mathematical models of a building's thermodynamics.
- **Exploration and exploitation:** RL thrives in stochastic environments. By continuously balancing the exploration of new actions with the exploitation of known rewarding

strategies, the agent can dynamically adapt to shifting weather patterns, fluctuating energy market signals, and changing thermal loads. This makes it uniquely suited for modern demand-response applications.

4.2 Control Architectures and Algorithm Selection

To navigate the competing objectives of minimizing operational costs while maximizing human comfort, control methodologies must be carefully structured. This involves both the overarching architectural framework and the specific algorithm deployed to execute the control policy.

4.2.1 The Evolution from General to Constrained Architectures

The management of multi-objective optimization in RL generally evolves through two foundational paradigms:

General Reinforcement Learning: In a General RL setup, competing objectives—such as energy savings versus thermal comfort—are managed by blending them into a single, unified scalar reward function [14, 15].

- **Operational Advantages:** The primary strength of General RL is its dynamic flexibility. By adjusting the weighting factors within the reward signal, the framework allows the agent to seamlessly shift its operational priorities in response to varying external market signals or environmental conditions.
- **Architectural Vulnerabilities:** The reliance on a unified scalar reward constitutes the framework's primary limitation. The tuning of penalty weights between competing objectives is highly sensitive and complex. A slight hyperparameter imbalance or an unforeseen environmental state can prompt the agent to aggressively prioritize economic savings at the severe expense of thermal comfort, or vice versa. Furthermore, General RL treats critical physical and operational limits merely as "soft" penalties rather than absolute rules, which introduces significant stability risks when deployed in physical building environments.

Constrained Reinforcement Learning: Because General RL often struggles to guarantee safe operation when faced with rigid real-world limitations, **Constrained RL** is utilized to address these critical shortcomings [16, 17]. Rather than relying solely on the sensitive tuning of penalty weights, Constrained RL reformulates the control problem as a Constrained Markov Decision Process (CMDP). In this architecture, critical operational boundaries—such as absolute indoor temperature limits or strict maximum power thresholds—are removed from the general reward pool and enforced as soft constraints. The agent is explicitly trained to optimize its primary economic objective strictly within a mathematically defined feasibility region. This transition from soft penalties to boundaries resolves the inherent vulnerabilities of general RL, ensuring robust, safe, and reliable real-world operation even during extreme grid events.

Note a constraint on the total power consumption can replicate fixed operating envelop signal coming from the DNSPs. During the training process, the RL agent will be penalised more if the power consumption go beyond the considered threshold. Therefore, the agent is progressively encouraged to reduce the frequency and magnitude of constraint violations.

4.2.2 Suitability of the Soft Actor-Critic (SAC) Algorithm for Building Control

Regardless of whether a general or constrained architecture is employed, the choice of the underlying learning algorithm is critical to the system's success. For smart building applications, the Soft Actor-Critic (SAC) framework has emerged as a state-of-the-art solution [10]. SAC is an off-policy algorithm particularly well-suited for the continuous control tasks inherent to HVAC management (e.g., modulating variable air volume dampers or adjusting continuous temperature setpoints).

Its primary advantage lies in its entropy-augmented objective function; SAC operates by maximizing both the expected return and the entropy (randomness) of its policy. This inherent promotion of entropy encourages robust environmental exploration and actively prevents the agent from converging prematurely on suboptimal control strategies. The efficacy of this algorithm is well-supported by our recent research work, environment-adaptive SAC frameworks have been successfully deployed to manage complex, continuous action spaces and balance comfort with energy efficiency across distinct building typologies, including both single-zone residential and commercial office environments [10]. Consequently, SAC demonstrates superior sample efficiency and performance stability across varied runs. Its ability to reliably manage the highly stochastic and unpredictable conditions typical of physical thermodynamics makes it a critical enabler for transitioning RL control from theoretical simulation to safe, operational reality [11].

5 Experimental Framework and Reward Mechanisms

Building upon the theoretical distinctions between general and constrained RL architectures outlined in Section 4, this study establishes a comprehensive evaluation framework to assess their practical implications under realistic operating conditions. To systematically investigate how different optimization objectives influence the balance among grid flexibility, occupant comfort, energy consumption, and operating cost, eight control strategies are formulated across the general and constrained RL paradigms. These strategies are constructed using three primary reward formulations. Historical spot-price data from the Australian Energy Market Operator (AEMO) were used to provide the time-varying electricity-price signals during RL-agent training.

At each control interval t, the reward functions are constructed from three components: the thermal-comfort reward $R_{\text{thermal}}(t)$, the energy-consumption reward $R_{\text{energy}}(t)$, and the operating-cost reward $R_{\text{cost}}(t)$. These components are normalized to comparable numerical ranges before being combined, ensuring that the weighting factors represent the intended control priorities rather than differences in the physical units or numerical magnitudes of the individual components.

Let $T_{\text{in}}(t)$ denote the indoor air temperature and let $T_{\text{low}}(t)$ and $T_{\text{high}}(t)$ denote the lower and upper thermal-comfort boundaries, respectively. A time-varying reference temperature $T_{\text{ref}}(t)$ is defined as

$$T_{\text{ref}}(t) = \begin{cases} \dfrac{T_{\text{high}}(t) + T_{\text{low}}(t)}{2}, & \text{during occupied periods,} \\ T_{\text{high}}(t), & \text{during unoccupied summer periods,} \\ T_{\text{low}}(t), & \text{during unoccupied winter periods.} \end{cases}$$

The thermal-comfort reward is calculated as

$$R_{\text{thermal}}(t) = \begin{cases} 2 - \dfrac{2}{1 + \exp(-|T_{\text{in}}(t) - T_{\text{ref}}(t)|)}, & T_{\text{in}}(t) \in \left[T_{\text{low}}(t), T_{\text{high}}(t)\right], \\ 1 - \dfrac{2}{1 + \exp(-|T_{\text{in}}(t) - T_{\text{ref}}(t)|)}, & \text{otherwise.} \end{cases}$$

Accordingly, a positive reward is assigned when the indoor temperature remains within the allowable comfort range, whereas a negative reward is imposed when a comfort violation occurs. The exponential form produces a smooth variation in reward as the indoor temperature approaches or moves away from the comfort range.

The HVAC energy consumption is normalized as

$$E(t) = \text{clip}\left(\frac{e(t) - e_{\min}}{e_{\max} - e_{\min}}, 0{,}1\right),$$

where $e(t)$ is the HVAC energy consumption at time t, and $e_{\min}$ and $e_{\max}$ are the minimum and maximum values used for normalization. The corresponding energy-consumption reward is defined as

$$R_{\text{energy}}(t) = \begin{cases} 1 - E(t), & T_{\text{in}}(t) \in [T_{\text{low}}(t), T_{\text{high}}(t)], \\ -E(t), & T_{\text{in}}(t) > T_{\text{high}}(t) \text{ in winter,} \\ & \text{or } T_{\text{in}}(t) < T_{\text{low}}(t) \text{ in summer,} \\ E(t) - 1, & T_{\text{in}}(t) < T_{\text{low}}(t) \text{ in winter,} \\ & \text{or } T_{\text{in}}(t) > T_{\text{high}}(t) \text{ in summer.} \end{cases}$$

This formulation rewards low energy consumption when thermal comfort is maintained. When the indoor temperature is outside the comfort range, it distinguishes between HVAC energy use that helps restore comfort and energy use that further aggravates the temperature violation.

The instantaneous electricity cost is calculated as

$$C(t) = p(t)e(t),$$

where $p(t)$ is the AEMO spot price at time t. When the HVAC output is measured as power rather than interval energy, the cost is calculated as

$$C(t) = p(t)P_{\text{HVAC}}(t)\Delta t,$$

where Δt is the control interval expressed in hours. For use in a weighted reward, the cost is normalized as

$$\hat{C}(t) = \text{clip}\left(\frac{C(t) - C_{\min}}{C_{\max} - C_{\min}}, 0,1\right),$$

and the operating-cost reward is defined as

$$R_{\text{cost}}(t) = -\hat{C}(t).$$

Price-Weighted Reward

The price-weighted reward dynamically balances thermal comfort and energy consumption using a price-dependent weighting factor $w_{\text{price}}(t) \in [0,1]$:

$$R(t) = [1 - w_{\text{price}}(t)]R_{\text{thermal}}(t) + w_{\text{price}}(t)R_{\text{energy}}(t).$$

The weighting factor is calculated from the normalized electricity price:

$$w_{\text{price}}(t) = \text{clip}\left(\frac{p(t) - p_{\min}}{p_{\max} - p_{\min}}, 0,1\right),$$

where $p_{\min}$ and $p_{\max}$ are obtained from the electricity-price data used during training. During low-price periods, the controller places greater emphasis on thermal comfort. As the electricity price increases, the contribution of $R_{\text{energy}}(t)$ becomes progressively larger, encouraging the controller to reduce HVAC energy consumption during expensive market periods.

Outdoor-Weighted Reward

The outdoor-weighted reward adjusts the balance between thermal comfort and operating cost according to the severity of the outdoor thermal conditions:

$$R(t) = w_{\text{outdoor}}(t)R_{\text{thermal}}(t) + [1 - w_{\text{outdoor}}(t)]R_{\text{cost}}(t).$$

The outdoor-temperature-dependent weight is defined as

$$w_{\text{outdoor}}(t) = \frac{1}{1 + \exp[-k(|T_{\text{out}}(t) - T_{\text{ref}}(t)| - \delta)]},$$

where $T_{\text{out}}(t)$ is the outdoor air temperature, k controls the steepness of the transition, and δ determines the outdoor-temperature deviation at which the reward begins to shift its emphasis toward thermal comfort. Under mild outdoor conditions, $w_{\text{outdoor}}(t)$ remains relatively small, allowing operating-cost reduction to receive greater emphasis. As the outdoor conditions become more extreme, $w_{\text{outdoor}}(t)$ approaches one, causing thermal comfort to dominate the reward and relegating operating cost to a secondary objective.

Cost-Only Reward

The cost-only reward is implemented exclusively within the constrained RL framework. In this formulation, thermal comfort is represented through explicit operational constraints rather than being incorporated as a weighted term in the scalar reward. The agent therefore focuses solely on minimizing electricity expenditure:

$$R(t) = -\hat{C}(t).$$

The constrained RL agent maximizes this reward by minimizing HVAC operating cost while satisfying the prescribed thermal-comfort constraints. This formulation avoids the need to manually select a relative weighting between comfort and cost and enables a direct evaluation of whether constrained RL can achieve economically efficient operation without violating the specified comfort requirements.

However, it is important to note that the constrained RL formulation employs soft constraints rather than absolute hard constraints. The agent is trained to maintain the expected cumulative thermal-comfort constraint cost below a predefined limit. Therefore, temporary comfort violations may still occur, particularly during exploration, training, or periods of extreme operating conditions. However, these violations increase the constraint cost and trigger a larger Lagrangian penalty, encouraging the policy to reduce subsequent violations. The resulting controller therefore minimizes operating cost while maintaining thermal-comfort violations within an acceptable long-term budget, rather than guaranteeing satisfaction of the comfort bounds at every individual timestep.

To comprehensively account for varying occupant requirements and evaluate the impact of the "solution space" on agent performance, the strategies were tested against two distinct daytime thermal boundary sets: **T1** (Wide/Mild: 20°C - 26°C) and **T2** (Narrow/Strict: 22°C - 24°C). Nighttime requirements for both scenarios were maintained at a broader 15°C - 30°C. Table 5 summarizes the experimental scenarios.

Table 5 Summary of Control Strategies and Experimental Scenarios

Reward Mechanism	RL Framework	Constraint Focus	Scenarios
1. Price-Weighted	General RL	None	T1 & T2
	Constrained RL	Energy Usage	T1 & T2
2. Outdoor-Weighted	General RL	None	T1 & T2
3. Cost-Only	Constrained RL	Thermal Comfort	T1 & T2

5.1 Performance and Flexibility Analysis

The proposed RL strategies were evaluated using flexibility maps that characterize the relationship between external operating conditions and the resulting HVAC response. The two external signals considered are outdoor air temperature and electricity price, while HVAC cooling power is used to represent the system response. Historical AEMO NSW1 spot-price data were used as the electricity-price signal throughout the RL simulations. Figure 13 illustrates an example 10-day segment of the electricity-price and outdoor-temperature signals at the 10-minute control resolution used in this study.

The RL agents were trained over an 80-day simulation period beginning on 1 January 2024, with training repeated for 10 episodes, and were subsequently evaluated over a separate 50-day unseen period. For the flexibility analysis, only observations corresponding to active building operating hours (8am-6pm) were retained, as these periods represent the times during which HVAC operation is actively subject to the daytime comfort requirements and demand-response objectives. This will create a total of 3000 samples over the unseen test period.

The retained evaluation data were then grouped according to both outdoor temperature and electricity price. Outdoor temperature was divided into three quantile-based categories, denoted Cool, Mild, and Hot, while electricity price was divided into Negative, Low, and High categories, with High prices defined as those above the 60th percentile. For each temperature–price combination, all observations falling within the corresponding bin were collected and the mean HVAC cooling power was calculated. The value reported in each heatmap cell therefore represents the conditional average HVAC power demand under that combination of environmental and market conditions, rather than cumulative energy consumption or the sum of power across the evaluation period. The N value shown in each cell indicates the number of active-hour observations contributing to the corresponding average.

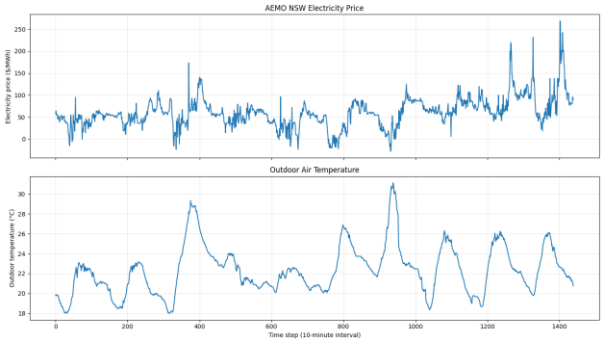


Figure 13 Representative 10-day external signals used in the RL evaluation: (a) raw AEMO NSW1 spot electricity price and (b) outdoor air temperature, both shown at a 10-minute control interval.

5.1.1 General RL under price-weighted reward

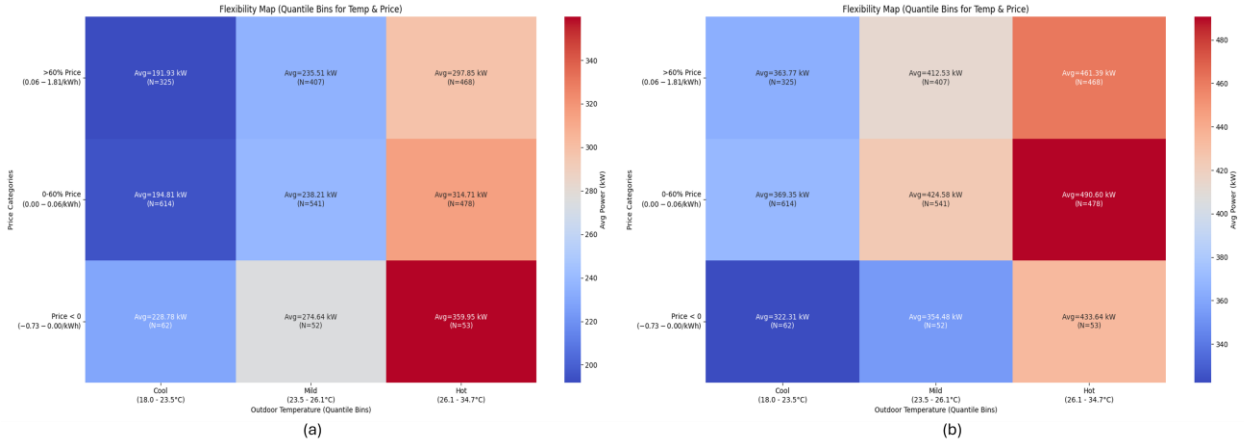


Figure 14 Performance of General RL under price-weighted reward: Comparison of (a) Wide Thermal Boundaries (T1) and (b) Narrow Thermal Boundaries (T2).

Figure 14 presents binned operating maps relating outdoor air temperature, electricity price, and HVAC power consumption. The maps were constructed using observations from the unseen test period during active building operating hours. At each timestep, the corresponding outdoor temperature, electricity price, and HVAC power consumption were recorded. The observations were then grouped into three outdoor-temperature categories (Cool, Mild, and Hot) and three electricity-price categories (Negative, Low, and High). The temperature categories were defined using tertiles of the outdoor-temperature distribution, while the price categories comprised non-positive prices, prices up to the 60th percentile, and prices above the 60th percentile. For each temperature–price bin, the mean HVAC power consumption of all observations falling within that bin was calculated and is reported in the corresponding heatmap cell; Ndenotes the number of observations contributing to that mean. Therefore, the maps do not represent a summation of electricity prices or power consumption across the test window, but rather the conditional average HVAC demand under different combinations of environmental and market conditions.

Figure 14(a) corresponds to the wide thermal-boundary scenario, T1 $(20–26^{\circ}\mathrm{C})$. The results show differences in power consumption across the price categories, although the response is not strictly monotonic across all temperature and price bins. In particular, the intermediate-price category exhibits higher consumption than the adjacent price categories under some outdoor-temperature conditions. This pattern should not be interpreted as the controller intentionally increasing consumption in response to intermediate electricity prices.

Two factors may contribute to this non-monotonic behaviour. First, the heatmap aggregates observations from different timesteps into broad quantile-based bins. Consequently, observations within the same price and temperature category may correspond to different indoor thermal states, cooling demands, and effects of previous control actions. Second, although the electricity price determines the weighting factor w_{price}, the resulting control incentive also depends on $R_{\mathrm{thermal}}$and R_{energy}. These reward components vary according to the current indoor temperature, HVAC energy consumption, and the indoor temperature's proximity to the applicable comfort boundaries. Therefore, the same electricity-price category may lead to different control priorities depending on the building's thermal state. Together with the thermal inertia of the building, these factors may explain the higher average consumption observed in the intermediate-price category.

Nevertheless, an endpoint comparison under Hot outdoor conditions $(26.1–34.7°C)$shows that average power consumption decreases from 359.95 kW during negative-price periods to 297.85 kW during high-price periods, defined as prices above the 60th percentile. This comparison suggests that the wider thermal range provides the controller with greater operational flexibility to reduce HVAC demand during expensive periods. However, the aggregated heatmap alone cannot determine whether this reduction was achieved through pre-cooling, immediate demand reduction, or another temporal control strategy.

Figure 14(b) presents the results under the narrow thermal-boundary scenario, T2 $(22–24°C)$, using the same RL framework and reward mechanism. Under these narrower boundaries, the relationship between electricity price and power consumption becomes less favourable and remains non-monotonic. Within the Hot temperature category, average power consumption increases from 433.64 kW during negative-price periods to 461.39 kW during high-price periods. This result indicates that the controller did not achieve an aggregate reduction in power consumption during high-price periods under the narrow thermal range.

A plausible explanation is that the narrower comfort range reduces the operational flexibility available for price-responsive control because the controller has less thermal slack within which HVAC operation can be shifted. However, because the heatmap aggregates observations from different timesteps and thermal states, the observed pattern cannot be attributed exclusively to the reward weighting or interpreted as direct evidence that the controller deliberately prioritizes thermal comfort over electricity price.

5.1.2 Constrained RL under price-weighted reward

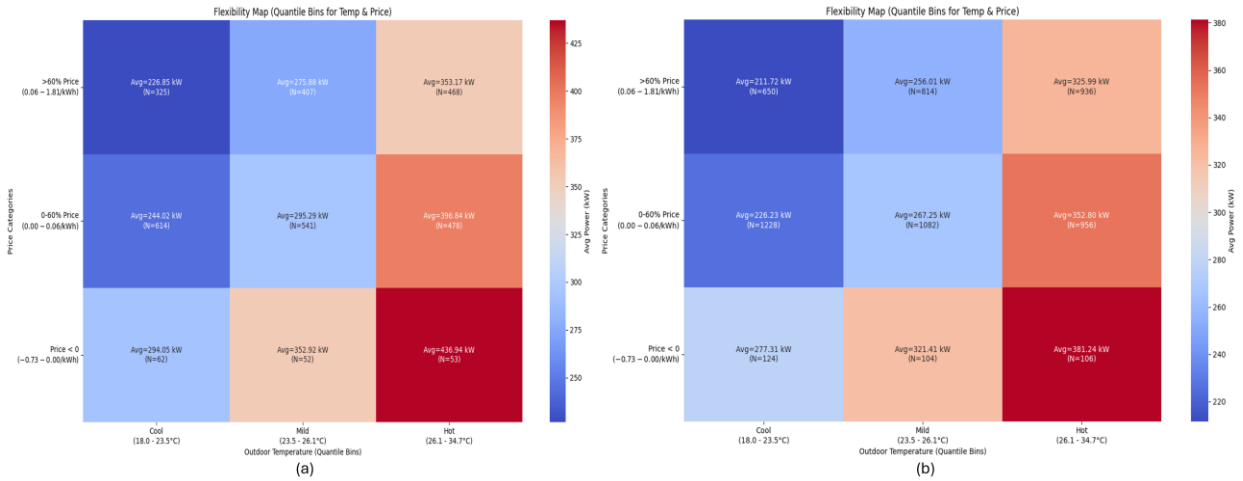


Figure 15 Performance of Constrained RL under price-weighted reward: Comparison of (a) Wide Thermal Boundaries (T1) and (b) Narrow Thermal Boundaries (T2).

Figure 15 presents the performance of the constrained RL framework. Unlike the general RL approach, which relies on a weighted scalar reward to balance thermal-comfort and energy-use objectives, the constrained RL agent incorporates an additional operational constraint cost associated with cooling demand. In this experiment, a 50 kW reference threshold was applied to the modelled cooling-demand signal used internally by the RL environment. No constraint cost is

incurred when this signal is at or below 50 kW, while values above the threshold incur an increasing constraint cost.

The 50 kW threshold should be interpreted as a soft optimization reference rather than a hard physical power limit. The constrained RL framework does not clip the control action or guarantee that cooling demand remains below 50 kW at every timestep. Instead, threshold exceedances generate a separate constraint cost whose contribution to policy optimization is regulated through an adaptive Lagrangian multiplier. In addition, the cooling-demand signal used to evaluate the constraint precedes the model bias correction used to reconstruct the reported cooling demand. Consequently, the bias-adjusted cooling-demand values shown in the operating maps can be substantially higher than the 50 kW internal reference threshold. The two quantities should therefore not be interpreted as directly equivalent pointwise power limits.

The constrained RL agent is optimized to reduce the expected accumulated constraint cost rather than to enforce a pointwise upper bound on cooling demand. When the allowable constraint budget is exceeded, the adaptive Lagrangian multiplier increases the optimization pressure associated with constraint violations, encouraging the agent to reduce their occurrence and severity over the optimization horizon. However, because the formulation is expectation-based and the instantaneous constraint cost is capped, it does not provide an absolute guarantee that the 50 kW reference threshold will never be exceeded. Comparing the average power consumptions of the constrained case in Figure 15 (b) with consumptions of the general RL in Figure 14(b) clearly reveals the reduction in power consumption, though it is not reaching the 50 kW target.

Figure 15 (a) shows the results under the wide thermal-boundary scenario, T1. Under Hot outdoor conditions, average cooling demand decreases from 436.94 kW during negative-price periods to 353.17 kW during high-price periods. This indicates that the constrained controller retains a clear aggregate price-responsive load-reduction capability when a relatively wide thermal operating range is available.

Under the narrow thermal-boundary scenario, T2, a similar price-responsive tendency remains visible. As shown in Figure 15 (b), average cooling demand under Hot conditions decreases from 381.24 kW during negative-price periods to 325.99 kW during high-price periods. In comparison with the corresponding general RL case, this suggests that the constrained formulation retains greater price-responsive flexibility when the allowable thermal operating range becomes more restrictive.

Overall, these results suggest that the constrained RL framework provides greater robustness to changes in the available thermal operating range, while maintaining an aggregate response to electricity-price signals. However, the operating maps aggregate observations from different timesteps and thermal states. They therefore demonstrate the resulting relationship between price, outdoor temperature and cooling demand, but do not by themselves identify the specific temporal control mechanism—such as pre-cooling or immediate load reduction—through which the response was achieved.

5.1.3 General RL under outdoor-weighted reward

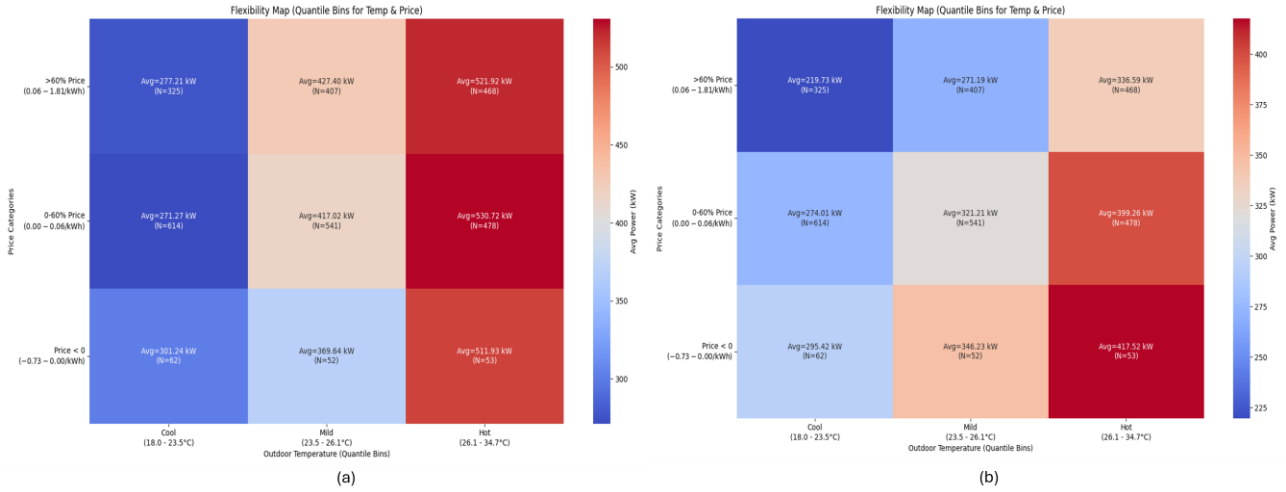


Figure 16 Performance of General RL under outdoor-weighted reward: Comparison of (a) Wide Thermal Boundaries (T1) and (b) Narrow Thermal Boundaries (T2).

Figure 16 illustrates the performance of the General RL agent under outdoor-weighted reward. In this mechanism, the weighting factor ($w_{outdoor}$) is determined by environmental severity rather than market signals. Consequently, the agent is designed to prioritize thermal comfort heavily during extreme outdoor conditions, relegating cost to a secondary objective.

This prioritization is clearly reflected in the flexibility maps, which show a distinct lack of price sensitivity. Unlike the vertical colour gradients seen in Figure 14 and Figure 15 (where power dropped as price rose), the columns in Figure 16 remain relatively uniform in colour regardless of the price level. In Figure 16 (a) (Wide Boundaries T1), the impact of the outdoor-driven weight is evident. Under 'Hot' outdoor conditions, the agent maintains extremely high-power consumption across all price bins to maximize comfort. During peak pricing events (>60% quantile), the average power consumption reaches 521.92 kW. This stands in sharp contrast to the Price-Driven Reward (Reward 1) shown in Figure 14 (a), which successfully reduced demand to 297.85 kW under identical conditions. This demonstrates that while Reward 2 leverages outdoor conditions to dynamically weight the trade-off between indoor thermal comfort and energy cost, economic flexibility is not prioritized, particularly when compared to other reward mechanisms within this test case.

A similar trend is observed in Figure 16 (b) (Narrow Boundaries T2), where the agent exhibits limited load-shedding capabilities in response to price increases. For instance, in the 'Hot' bin, consumption decreases only moderately from 417.52 kW (Negative Price) to 336.59 kW (High Price).

5.1.4 Constrained RL under cost-only reward

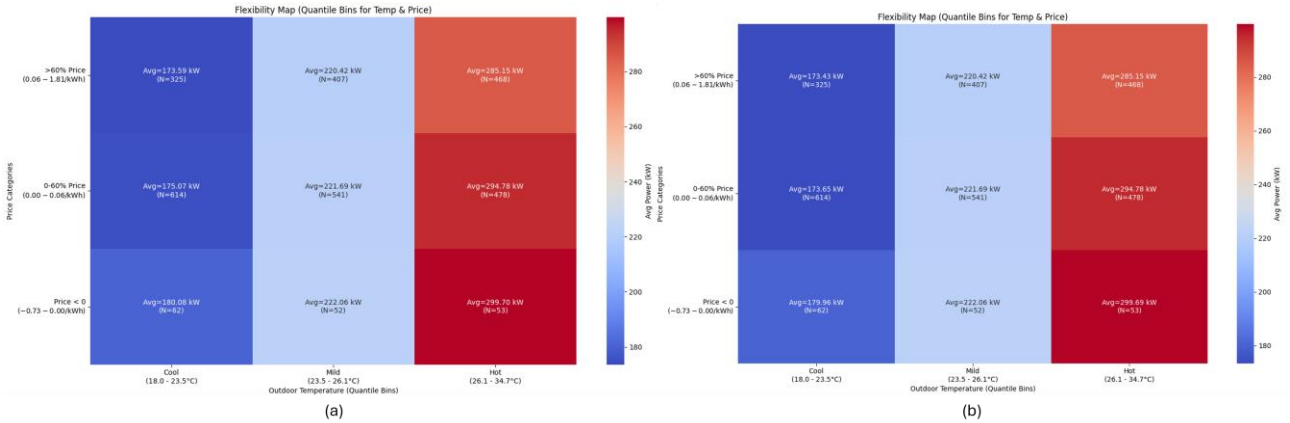


Figure 17 Performance of Constrained RL under cost-only reward: Comparison of (a) Wide Thermal Boundaries (T1) and (b) Narrow Thermal Boundaries (T2).

Figure 17 presents the performance of the Cost-Only reward mechanism implemented within the Constrained RL framework. In this formulation, the agent's primary objective is strictly to minimize operational costs, while thermal comfort is handled exclusively as a constraint. The most notable observation in this scenario is the overall reduction in energy intensity. Across all temperature and price bins, this strategy yields the lowest power consumption among all tested approaches. For example, under 'Hot' conditions and peak pricing, the average consumption is approximately 285.15 kW, significantly lower than the 325 kW observed in Reward 1 (Constrained) or the 336 kW in Reward 2. This indicates that a pure cost-minimization objective drives the agent to operate at the lower bound of the feasibility region, consuming only the bare minimum energy required to prevent constraint violations.

Furthermore, comparing Figure 17 (a) (Wide Boundaries T1) and Figure 17 (b) (Narrow Boundaries T2) reveals a remarkably consistent performance profile. This suggests that the "Cost-Only" signal is so dominant that the agent does not utilize the additional "thermal slack" provided by the wide T1 boundaries to improve comfort. Instead, regardless of whether the boundary is wide or narrow, the agent converges on the most efficient operational state allowed, demonstrating high stability but potentially lower comfort flexibility compared to the weighted reward mechanisms.

5.1.5 Aggregated Performance of All Strategies

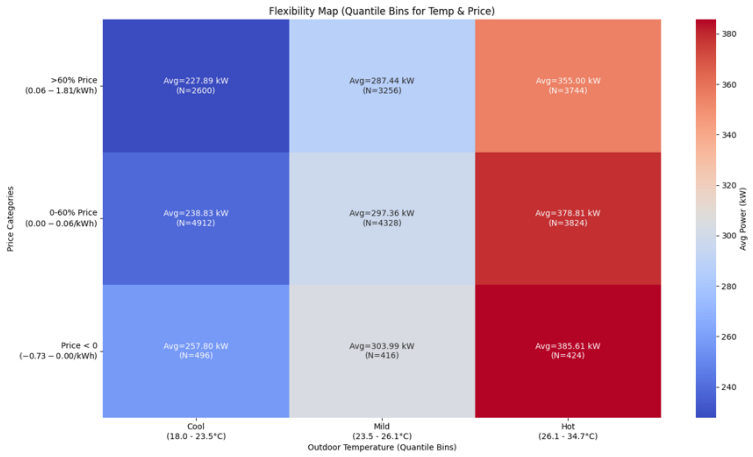


Figure 18 Aggregated Flexibility Map: Average performance across all 8 tested control strategies and thermal boundary conditions

Figure 18 presents the comprehensive flexibility map, aggregating the data from all 8 control strategies and both thermal boundary scenarios (T1 and T2). This visualization provides a holistic view of the "systemic flexibility potential" of the smart building environment when managed by the proposed portfolio of RL agents. The aggregated results reveal a robust, albeit smoothed, correlation between market signals and power consumption, where a consistent load-shifting trend remains visible. As electricity prices rise (moving from bottom to top), the average power consumption decreases across all temperature bins.

Most notably, in the 'Hot' outdoor condition, the aggregated system reduces its average power consumption from 385.61 kW during negative pricing events to 355.00 kW during peak pricing events. This represents a net reduction of approximately 30 kW. This reduction confirms that, on average, the proposed RL frameworks can deliver grid flexibility without compromising the overall operational stability of the building.

5.2 Quantitative Assessment via the Flexibility Index (FI)

To rigorously quantify the demand response capabilities of the proposed control architectures, we evaluated their performance using a formulated Flexibility Index (FI), adapted from the methodology proposed by Junker et al. for characterizing building energy flexibility [12].

The FI measures the capacity of a system to shift its operational behaviour in response to an external penalty signal—in this context, real-time electricity prices. The index compares the performance of a "penalty-aware" control strategy against a "penalty-ignorant" baseline. For this study, the baseline was established using a published general RL- based benchmark agent [10]. it trained exclusively on an outdoor-driven reward mechanism, meaning it lacked any visibility of market price signals [10]. On the other hand, all other approaches are trained with price signal.

The Flexibility Index based on real-time operational costs is defined as:

$$FI_{cost} = 1 - \frac{\sum Price(t) \times Energy_{w\ flexibility}(t)}{\sum Price(t) \times Energy_{wo\ flexibility}(t)}$$

To distinguish the economic benefits of price-responsive load shifting—including demand-side arbitrage through pre-cooling or pre-heating—from physical reductions in peak-period demand, the index was also evaluated using absolute HVAC energy consumption during predefined high-price periods:

$$FI_{energy} = 1 - \frac{\sum Energy_{w\ flexibility}(t)}{\sum Energy_{wo\ flexibility}(t)}$$

A positive FI_{cost} indicates that the price-aware controller achieves lower operating costs than the price-ignorant baseline, including savings obtained through temporal demand shifting and price arbitrage. A positive FI_{energy} indicates that the controller also reduces HVAC energy consumption during predefined high-price periods. Negative values indicate that the price-aware strategy performs worse than the baseline under the corresponding metric, potentially due to inefficient pre-conditioning, restrictive thermal requirements, or conflicts between optimization objectives.

The calculated Flexibility Indices across the wide (T1: 20°C - 26°C) and narrow (T2: 22°C - 24°C) thermal boundaries reveal critical insights into the robustness of the evaluated architectures.

The Vulnerability of General RL (Reward No. 1)

Under the wide T1 boundary, the Price-Weighted General RL approach (R1) managed a marginal positive flexibility (FI_cost ∼ 0.37%), utilizing the broad "thermal slack" to navigate price signals. However, when subjected to the strict T2 boundary, the performance of this architecture collapsed drastically, plummeting to an FI_cost of -62.56% and an FI_energy of -68.99%. This extreme degradation quantitatively proves the primary vulnerability of General RL: when the physical solution space tightens, the agent abandons economic optimization entirely to avoid massive thermal discomfort penalties, resulting in severely rigid and inefficient grid behaviour.

The Superior Robustness of Constrained RL (Reward No. 3)

In stark contrast to the General RL models, the Cost-Only constrained RL approach (R3) demonstrated unparalleled robustness. By treating thermal limits as absolute, mathematically defined constraints rather than weighted penalties, the agent's sole objective was to minimize the cost penalty within the feasible region.

As a result, the R3 strategy was the only approach to yield consistently strong, positive Flexibility Indices regardless of the thermal boundary width. Under T1, it achieved an FI_cost of 6.24% and an FI_energy of 6.41%. Remarkably, under the highly restrictive T2 scenario, it improved its relative performance, achieving the highest overall flexibility scores (FI_cost = 7.51%, FI_energy = 7.63%).

5.3 Price-demand mapping

Creating a function that links regional price to nodal demand is the intention of this section. To achieve this aim, we build a sequence-to-sequence regression model that learns the mapping from energy price, outdoor temperature, and hour-of-day to building electrical load (or cooling usage). The main idea is to predict not a single scalar output, but an entire window of daily load values given a corresponding window of input features. A one-to-one regression case is also investigated. In this setting, each sample contains one value of each input feature, namely energy price, outdoor temperature, and hour of day, and the model predicts the corresponding single value of building electrical load or cooling usage. This differs from the sequence-to-sequence setting, where a 24-timestep window of input features is mapped to a 24-timestep target sequence. The one-to-one model estimates the instantaneous relationship between the input variables and demand.

The data used for creating such a mapping function come from RL simulation outputs. Each input data file represents a different RL agent configuration, the description of which can be found in Table 5. To prepare the input data files, we slice out only the unseen period of each simulation to the RL (rows beyond a fixed index) and further filter to active occupancy hours by removing rows where the comfort boundary indicates unoccupied conditions. An aggregated version of all the cases is also assessed, in which the target column is the element-wise sum of load for all simulation files. Features in this case come from a single reference file, making the implicit assumption that exogenous conditions (price, temperature, time) are shared among scenarios.

For the regression task, we create sliding-window samples from the time series. Given a window length of 24 timesteps, each training example consists of the concatenation of all feature sequences over that window (so for 3 features × 24 steps = 72 input dimensions) mapped to the 24-step target sequence. This transforms the problem from a standard tabular regression into a multi-output regression in which the model must predict all 24 steps simultaneously.

The train/test split is done per file using an 80/20 chronological split (no shuffling to preserve temporal order), and the resulting windows from all files are then concatenated into global train and test arrays.

The model itself is a pipeline of standard scaling followed by a Ridge (linear) regression. Because the target is multi-dimensional (24 outputs), scikit-learn's Ridge natively fits a separate linear equation for each output step, but with shared regularisation. The scaler normalises the concatenated feature vectors to zero mean and unit variance before the linear model sees them.

Cross-validation uses a time-series aware splitting strategy that respects temporal ordering (each fold's validation set always comes strictly after its training set). The metrics computed both on the full flattened predictions (all steps pooled together) and per sequence step to check if prediction quality degrades at later horizons. The following tables 6-10 summarises the regression metrics for each of the 8 different settings in Table 5, as well as an aggregated ensemble of all of them. Additionally, Figure 19 shows the performance of the one to one and sequence to sequence models for predicting cooling usage and total load on test and train data. The developed models follow the following general mathematical formula:

$$M_c(t) = A_c P(t) + B_c T(t) + C_c H(t) + D_c$$

$$M_t(t) = A_t P(t) + B_t T(t) + C_t H(t) + D_t$$

Where $M_c(t)$ and $M_t(t)$ are cooling usage and total demand, respectively, $P(t)$ is the price $T(t)$ is the outside air temperature, $H(t)$ is time of day. The regression constants A, B, C are slopes of the learned linear model and D is its bias term. The slopes and bias terms of the derived models have been provided in Table 11 to Table 14.

The comparison between the one to one and sequence to sequence settings indicates that using a 24 timestep input window provides a small but consistent improvement. For cooling usage, the mean test R-squared increases only slightly from approximately 0.80 to 0.82. This suggests that much of the cooling response can be explained by the instantaneous values of price, outdoor air temperature, and hour of day. The sequence-to-sequence model still reduces the test RMSE in most cases, specifically for the aggregated case. This indicates that short-term temporal context contains useful information, although its added value is limited when the target variable is already strongly coupled with the current operating conditions.

The cooling demand is more directly governed by the selected explanatory variables, especially outdoor temperature and time of day. In contrast, total site load likely contains additional components that are not represented in the input feature set, such as plug loads, lighting, equipment schedules, stochastic occupant behaviour, and other non-HVAC demands. This is evidenced based on better metrics obtained for predicting cooling usage compared to that of total site load. The sequence-to-sequence setting improves total load prediction more clearly than cooling prediction, but the remaining error indicates that price, temperature, and hour of day alone might be insufficient to fully represent whole building demand dynamics. Therefore, the obtained mapping is more reliable for estimating price-responsive cooling behaviour than for explaining total site electricity use.

To assess if the modelled response to electricity price represented a meaningful signal and is not an effect arising by chance, a day-block permutation test was performed. First, the contribution of price to the fitted model's prediction of chiller load was quantified. The price series was then permuted by whole days 500 times, while preserving the intraday price profile and hour of day structure while breaking the association between price and load. For each permutation, the same modelling pipeline was refitted and the price effect was recalculated. These permutation results formed a null distribution representing the magnitude of apparent price response expected under the null hypothesis that price provides no additional information about load beyond temperature and hour of day. The observed price response was then compared with this null distribution, and permutation-based p-values were calculated. An observed response substantially greater than that obtained under the null distribution was considered statistically significant.

Table 6 Permutation-test results for the significance and detectability of the modelled price response

Case	Significant cases	Median observed response (% of load)	Median detection threshold (null 95th percentile)
1-to-1 – cooling_usage	7/9	0.79%	0.41%

seq24 – cooling_usage	7/9	1.55%	1.13%
1-to-1 – total_load	1/9	0.61%	1.07%
seq24 – total_load	0/9	1.32%	2.69%

The results indicate that the price response is generally distinguishable from the permutation null for cooling usage. In both the 1-to-1 and seq24 configurations, 7 of the 9 models showed statistically significant price responses, suggesting that the relationship between electricity price and cooling demand can be recovered by the models beyond the effects of temperature and hour of day. The median observed price response was also greater than the corresponding median null detection threshold in both cases. The exceptions were Constrained_RL_R3_T1 and Constrained_RL_R3_T2, which exhibited near-zero price responses. In contrast, the price response was not reliably detectable for total load. Only 1 of 9 models was significant in the 1-to-1 configuration, while none were significant in the seq24 configuration. In both total load cases, the median observed response was below the null detection threshold. This is consistent with the expectation that the price-driven cooling response becomes diluted when cooling demand is aggregated with whole building electricity consumption.

Table 7 Results for 1-1 regression between price, outside air temperature and hour of day and site cooling usage

Model	train_r2	test_r2	train_rmse	test_rmse	train_mae	test_mae
General_RL_R1_T1	0.89	0.90	18.72	18.36	9.60	10.85
General_RL_R1_T2	0.37	0.45	66.89	71.59	50.56	57.07
General_RL_R2_T1	0.87	0.88	41.41	38.92	34.93	32.00
General_RL_R2_T2	0.39	0.40	71.27	78.52	60.47	68.62
Constrained_RL_R1_T1	0.84	0.83	29.03	31.84	22.02	24.52
Constrained_RL_R1_T2	0.81	0.84	21.77	20.92	11.83	13.76
Constrained_RL_R3_T1	0.99	1.00	3.81	3.89	3.25	3.26
Constrained_RL_R3_T2	1.00	1.00	3.42	3.87	2.91	3.19
Summed_All_Files	0.95	0.94	111.90	124.81	85.24	97.66

Table 8 Results for 1-1 regression between price, outside air temperature and hour of day and site total load.

Model	train_r2	test_r2	train_rmse	test_rmse	train_mae	test_mae
General_RL_R1_T1	0.37	0.55	62.15	50.80	50.64	38.00
General_RL_R1_T2	0.17	0.37	97.25	86.40	80.94	70.43
General_RL_R2_T1	0.61	0.67	78.26	76.30	65.04	60.30
General_RL_R2_T2	0.21	0.35	91.88	88.55	73.09	72.90
Constrained_RL_R1_T1	0.45	0.65	64.43	53.91	52.34	41.42
Constrained_RL_R1_T2	0.28	0.51	63.87	48.17	51.84	35.73
Constrained_RL_R3_T1	0.37	0.57	61.89	51.24	51.32	39.27
Constrained_RL_R3_T2	0.37	0.57	61.62	51.09	51.03	39.25
Summed_All_Files	0.41	0.62	504.33	422.39	419.86	325.76

Table 9 Results for sequence to sequence (24 of each input to 24 of target) regression between price, outside air temperature and hour of day and site cooling usage.

Model	train_r2	test_r2	train_rmse	test_rmse	train_mae	test_mae
General_RL_R1_T1	0.90	0.90	17.95	18.10	9.94	11.12
General_RL_R1_T2	0.46	0.45	62.17	72.42	47.15	58.36
General_RL_R2_T1	0.88	0.89	39.72	36.98	33.48	30.09
General_RL_R2_T2	0.45	0.44	67.47	76.21	56.32	66.44
Constrained_RL_R1_T1	0.88	0.87	25.29	28.19	19.50	21.87
Constrained_RL_R1_T2	0.86	0.85	19.14	20.15	11.49	13.51
Constrained_RL_R3_T1	0.99	1.00	3.72	4.02	3.10	3.34
Constrained_RL_R3_T2	1.00	1.00	3.34	3.94	2.81	3.27
Summed_All_Files	0.96	0.95	100.00	116.86	77.73	93.70

Table 10 Results for sequence to sequence (24 of each input to 24 of target) regression between price, outside air temperature and hour of day and site total load

Model	train_r2	test_r2	train_rmse	test_rmse	train_mae	test_mae
General_RL_R1_T1	0.42	0.59	59.47	48.85	47.25	34.75
General_RL_R1_T2	0.28	0.41	90.71	83.20	75.00	67.24
General_RL_R2_T1	0.65	0.70	73.46	74.03	61.48	57.73
General_RL_R2_T2	0.28	0.38	87.72	87.19	68.65	71.11
Constrained_RL_R1_T1	0.51	0.68	61.33	51.99	49.00	38.51
Constrained_RL_R1_T2	0.35	0.55	60.59	46.17	48.23	32.67
Constrained_RL_R3_T1	0.44	0.61	58.40	48.88	47.64	35.16
Constrained_RL_R3_T2	0.44	0.61	58.16	48.71	47.38	35.18
Summed_All_Files	0.47	0.65	477.78	408.63	391.87	299.42

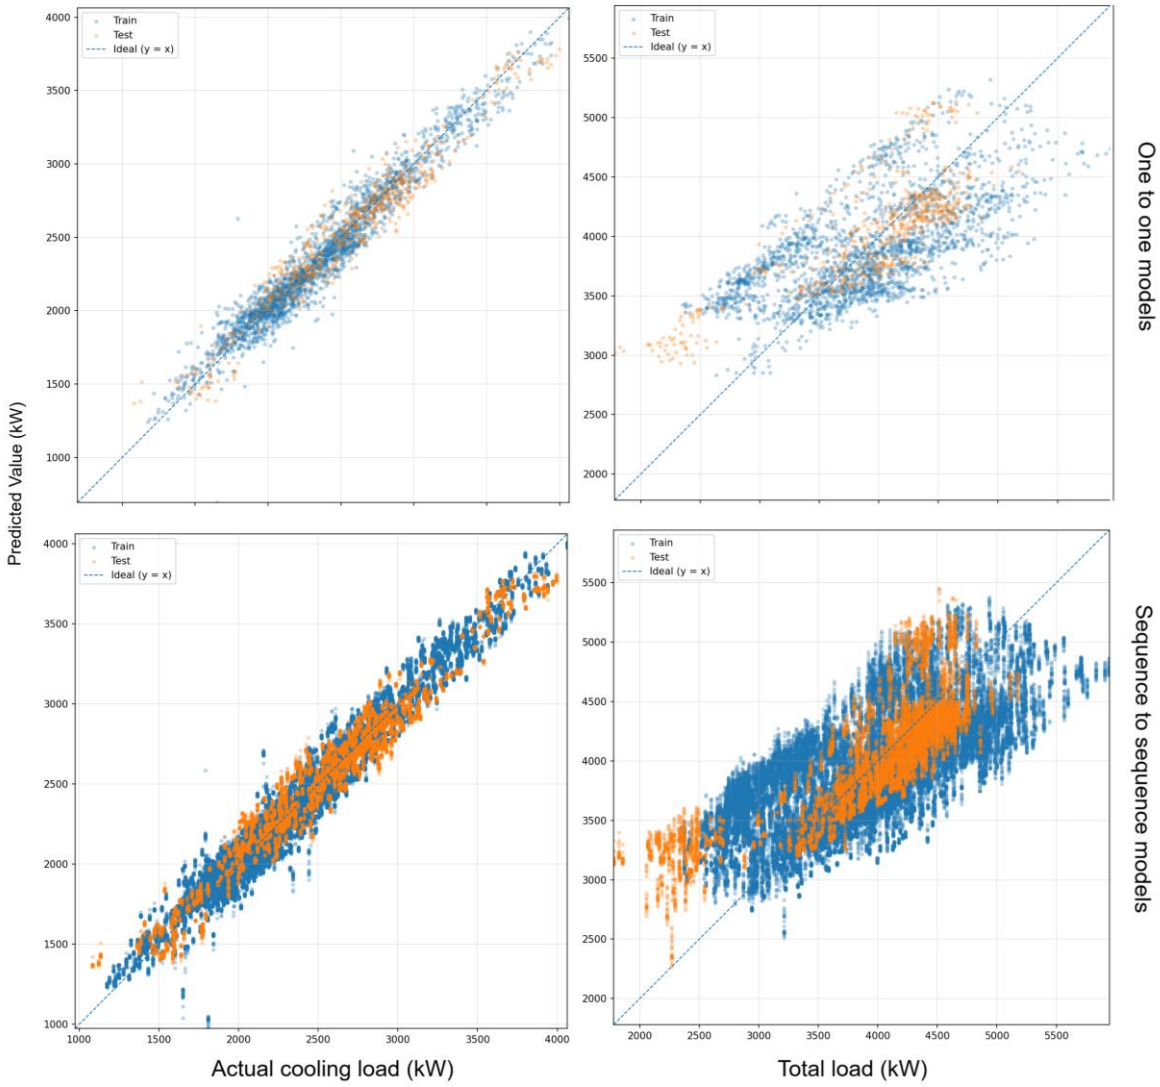


Figure 19 Performance of the developed sequence to sequence and one to one models for predicting total load and cooling usage. The top plots show the one-to-one models while the bottom plots show results for sequence to sequence models.

Table 11 Slopes and bias terms for the one to one cooling usage models.

Model	slope_price A_c	slope_outdoor_temp B_c	slope_hour C_c	Bias D_c
General_RL_R1_T1	-0.05518	18.05032	0.011514	-202.313
General_RL_R1_T2	-0.18938	17.19595	-0.49996	4.952389
General_RL_R2_T1	0.074839	36.15931	0.064647	-507.278
General_RL_R2_T2	-0.28182	18.75392	-2.02784	-119.62
Constrained_RL_R1_T1	-0.19025	22.36973	-0.48203	-241.944
Constrained_RL_R1_T2	-0.06372	15.52899	-0.80218	-125.398
Constrained_RL_R3_T1	-0.00124	17.91742	-0.27228	-217.654
Constrained_RL_R3_T2	0.000188	17.96409	-0.28541	-219.124
Summed_All_Files	-0.70637	163.8651	-4.28544	-1626.63

Table 12 Slopes and bias terms for the one to one total load models.

Model	slope_price A_t	slope_outdoor_temp B_t	slope_hour C_t	Bias D_t
General_RL_R1_T1	0.022076	16.55122	0.071135	16.29881
General_RL_R1_T2	-0.09028	15.74882	-1.10556	227.0936
General_RL_R2_T1	0.167743	33.41422	0.145567	-257.548
General_RL_R2_T2	-0.19289	16.48628	-1.52609	112.9148
Constrained_RL_R1_T1	-0.1144	20.72676	-0.34494	-20.2024
Constrained_RL_R1_T2	0.01486	13.92984	-0.60464	94.14533
Constrained_RL_R3_T1	0.083528	16.22024	-0.26085	5.730932
Constrained_RL_R3_T2	0.084527	16.37376	-0.28949	1.751004
Summed_All_Files	-0.02484	149.4511	-3.91487	180.1843

Table 13 Slopes and bias terms for the sequence to sequence cooling usage models.

Model	slope_price A_c	slope_outdoor_temp B_c	slope_hour C_c	Bias D_c
General_RL_R1_T1	-0.00571	0.76965	-0.01522	-203.145
General_RL_R1_T2	-0.01322	0.707249	-0.04334	24.81365
General_RL_R2_T1	0.009898	1.519185	0.005638	-525.418
General_RL_R2_T2	-0.02541	0.801452	-0.07034	-116.28
Constrained_RL_R1_T1	-0.0183	0.954865	-0.00717	-244.45
Constrained_RL_R1_T2	-0.00643	0.660432	-0.00113	-137.643
Constrained_RL_R3_T1	-3.13E-05	0.746515	-0.01874	-215.474
Constrained_RL_R3_T2	0.000121	0.752075	-0.02002	-218.994
Summed_All_Files	-0.05909	6.911424	-0.17032	-1636.59

Table 14 Slopes and bias terms for the sequence-to-sequence total load models.

Model	slope_price A_t	slope_outdoor_temp B_t	slope_hour C_t	Bias D_t
General_RL_R1_T1	0.007776	0.596479	-0.04367	76.96319
General_RL_R1_T2	0.002069	0.535361	-0.04351	291.556
General_RL_R2_T1	0.024734	1.27898	0.020102	-218.145
General_RL_R2_T2	-0.01083	0.600332	-0.06376	169.2696
Constrained_RL_R1_T1	-0.00476	0.777984	-0.0373	38.44794
Constrained_RL_R1_T2	0.007187	0.480775	-0.0286	146.1257
Constrained_RL_R3_T1	0.014045	0.564123	-0.03148	64.5475
Constrained_RL_R3_T2	0.014153	0.571462	-0.03369	60.23696
Summed_All_Files	0.054374	5.405495	-0.2619	629.0019

Part III Whole system simulation

6 Impact of price-responsive resources on dispatched price

For a whole system simulation from market-level optimisation (see part I) to local building level optimisation (see Part II), the SNEM 2000 grid model explained in section 2.2 should be coupled with HVAC Gyms described in section 3. This would create a closed loop for updating substation level demand using Gym and retrieving the resultant regional prices from the NEM model.

6.1 General network transformations

The optimisation framework in Part I, co-optimises the electricity and FCAS market at a given snapshot. For the closed-loop simulation of the whole system, however, a timeseries simulation is required. It basically means all the network parameters and inputs, including all the generator energy and FCAS bids, need to be updated at each iteration. Since such information are not available, we made the following simplifying assumptions:

- We take all the active power load values at all buses in the existing sNEM 2000 snapshot model and scaling them according to a given daily load profile such that the existing snapshot values occur at the time of daily peak demand.
    - We do the reactive power load scaling by preserving the original power factors of all loads across all time stamps.
    - In our experiments and to calculate the scaling factors, we use the NEM total demand across a typical week in February 2025. We first find the maximum demand across the entire week and then dividing all the other demand values by the maximum value to find the scaling factors over the entire week. The proposed scaling factors and the resulting dispatch price can be seen in Figure 20 with undesirable high pricing of up to $40,000 per MW.

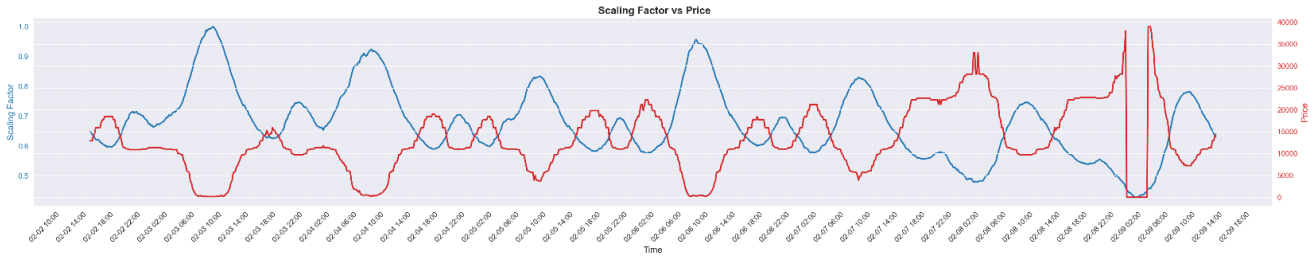


Figure 20 Scaling factors for all the loads in NEMX.jl and the resulting calculated dispatch price, when we run NEMX in timeseries mode.

    - Note our proposed scaling method is a simple way to extend the existing valid snapshot load values across the NEM. While there might be better ways to scale the loads for time series analysis, guaranteeing the feasibility of optimisation at each time step is not an easy task and should be considered as part of any future scaling method improvement. Ultimately, an ideal scaling method should replicate NEM

price signals and be validated against them. Our proposed approach is just for the proof of concept and has not been validated against NEM prices.

- As it can be seen in Figure 20, the dispatch price is inversely related to demand, which is aligned with bidding curves of scheduled loads in Figure 21. In this illustrative example, only band 9 and 10 capacity will be dispatched as their price bids are above the spot price. In other words, the remaining bands won't be dispatched, since the spot price is not low enough to justify consumption on those bands.

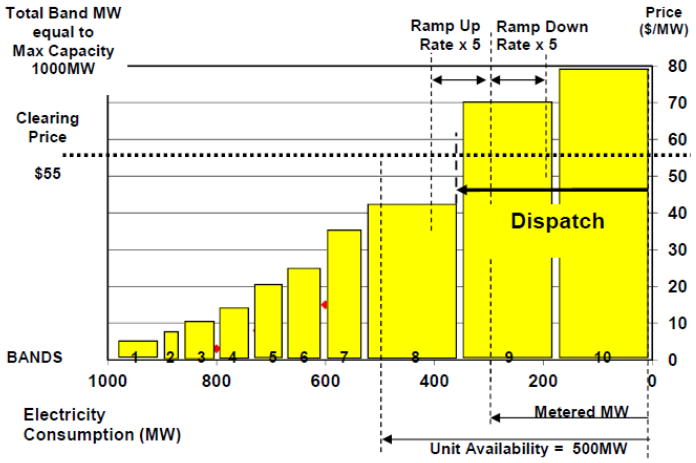


Figure 21 Price-demand bids of scheduled loads in NEM, taken from AEMO's guide to scheduled loads, July 2021

- The very high prices in Figure 20 in low demand hours is likely because of lack of tuning in the generators and scheduled demands bidding curves. Since updating ~300 of such curves is almost impossible, we propose a simple solution by using the same well-behaved arbitrary price-capacity bidding curve for all generators in the NEM. This selection ensures the flexible load demand across a day is the only deriving factor for changes in the price.
    - In our simulation, we used the following 10 quantity-price bidding values (the odd elements of this vector are the capacity, and the even elements are the associated price)

```
gen["cost"] = [0.0, -967.3, 1.1, -967.3, 4.375, 18.22, 8.75, 28.22, 8.75,
38.22, 7.75, 48.22, 7.75, 58.22, 7.75, 68.22, 7.75, 18.22, 0.0, 16057.2]
```

    - Note this well-behaved arbitrary price-capacity curve will remain the same across the entire window of a timeseries simulation.
    - Later in our experiments, we updated our quantity-price bidding values based on the fuel type of generation, but still the same across different time stamps.

```
if gen["fuel"] == "Coal"
            gen["cost"] = [0.0 0.0 2 0.98 4.0 9.85 6.0 35.58 6.5 54.02 7.0
83.74 7.5 137.9 8.0 216.69 9.0 315.14 10.0 1634.768]
        elseif gen["fuel"] == "Water"
            gen["cost"] = [0.0 0.0 1.5 0.0 4.0 10.07 6.0 31.83 6.5 49.06 7.0
72.63 7.5 90.77 9.0 117.94 10.0 281.2 11.0 1505.288]
        elseif gen["fuel"] == "Wind"
```

```
            gen["cost"] = [0.0 -965.5 0.5 -964.63 1.0 -182.5 1.5 -84.75 2.0 -
68.85 2.5 -45.05 3.0 -34.73 3.5 -2.99 6.0 19.51 7.0 160.96]
        elseif gen["fuel"] == "Solar" || gen["fuel"] == "Biomass" ||
gen["fuel"] == "CapBank/SVC/StatCom/SynCon" || gen["fuel"] == "Battery"
            gen["cost"] = [0.0 -984.7 0.1 -738.53 0.5 -172.32 1.0 -64.01 1.5
-58.1 3 -54.16 4.0 32 4.5 64 5.0 128 8.0 163.02]
        elseif gen["fuel"] == "Natural Gas" || gen["fuel"] == "Distillate"
            gen["cost"] = [0.0 0 1.0 0 2.0 124.22 2.5 198.95 3.0 298.11 3.5
377.6 4.0 557.49 4.5 1366.25 5.0 12841.65 8.0 16496.2]
end
```

- It can also be seen in Figure 20 than when the scaling factors are below 0.47, the market optimisation problem becomes infeasible and the price drops to zero. To avoid that, we consider a lower cap of 0.47, if the scaling factor becomes lower than that.

Remark. After the three above mentioned adjustments to the NEMX package, we can run a timeseries co-optimisation problem with reasonable price values as it can be seen in the following sections. It should be noted that in this simple modification, the only variable input at each time step is the outside air temperature and demand values of flexible temperature-dependant loads and everything else including the generator/demand bids remain the same as in the snapshot scenario. While not ideal, this is just for the proof of concept. Considering all the other variable factors at each time step means building the NEM dispatch engine, which has never been the intention of this project. It should also be noted that since we put our flexible demand in the NSW region, we only show the NSW regional reference price at bus 130.

6.2 Connecting a flexible demand to sNEM2000 model

We consider Newcastle Energy Centre chiller demand as an example of flexible demand in this study. However, since the total demand of the Newcastle site is around 1MW, see Figure 22, it is not yet at the required aggregated level to be connected to a (sub transmission) node in the sNEM 2000 model. Therefore, we have considered a constant total demand multiplying factor that simply multiplies the current demand of the Newcastle site by that factor before its connection to the sNEM 2000 model.

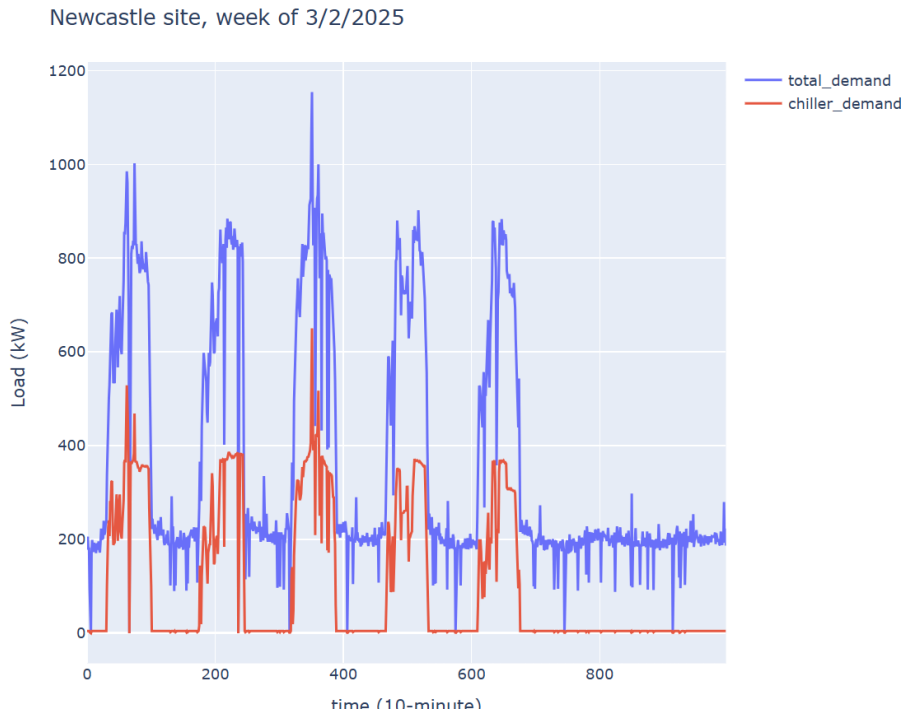


Figure 22 Flexible chiller demand and total demand of Newcastle Energy Centre in a typical week in February 2025

Since the physical location of nodes in the synthetic NEM model is not linked to the exact location of loads in the NEM, we need to manually link a node to Newcastle based on some intuitions. Based on the type of transmission lines and generator capacities/types, illustrated in Figure 23, we have identified the Node 690 of the sNEM 2000 model as the closest representation of the Newcastle site flexible demand.

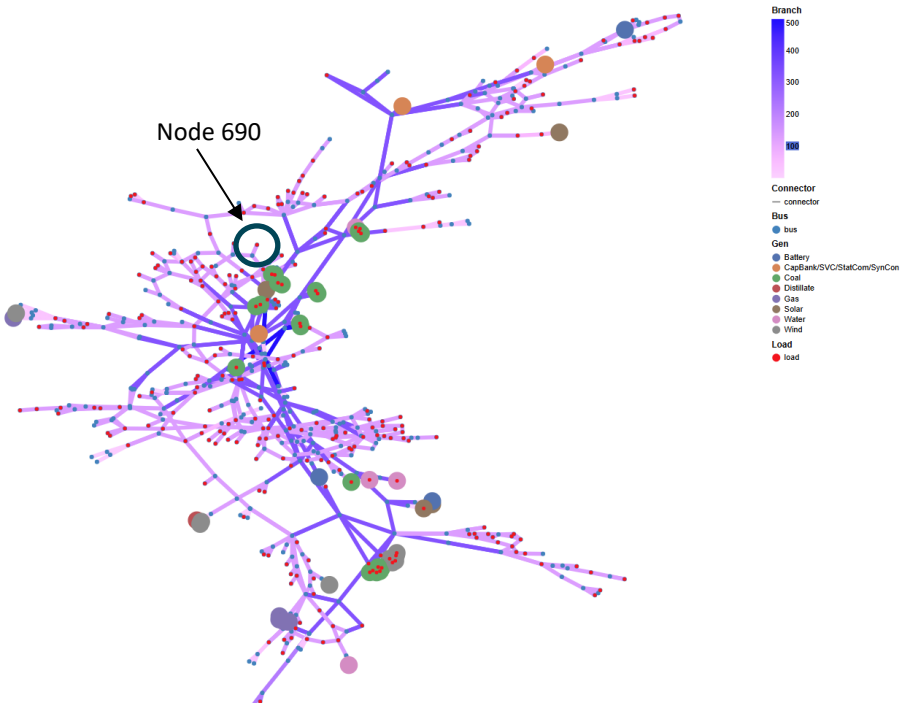


Figure 23 NSW section of the sNEM 2000 model and node 690 as Newcastle site representation

In the closed loop simulation, the specified node 690 active power is overwritten by the summation of

- underlying known non-chiller demand (fixed at each time step), and
- the dynamic flexible chiller power (HVAC Gym output varying at each step),

multiplied by the total demand multiplying factor.

6.3 Python-Julia wrapper function

For the closed loop simulation of the whole system, a wrapper function was developed in Python to set up and solve a Julia-based DC optimal power flow with co-optimised FCAS for one point in time (10-minute temporal resolution herein) during the desired simulation period. Each call to the function receives a new flexible demand value from the Newcastle Gym, a timestamp string, and a total demand multiplying factor.

The wrapper begins by looking up the corresponding timestamp in a static data frame containing Newcastle site's underlying (non-flexible) load and NEM total demand time-series. The underlying load will be added to the flexible demand value (output of the Gym after local RL price-based optimisation) to form the Newcastle site total demand. This total demand will then be multiplied by the constant multiplying factor before its attachment to the sNEM2000 model and running a market optimisation to find the price of the next step. The new price will ultimately change the flexible portion of the demand due to the local price-based optimisation, and the closed loop continues.

6.4 Simulation results

We first investigate the impact of non-scheduled flexible loads on dispatch price over the first day of a typical week shown in Figure 22. Simulation time step is 10 minute. For simplicity, we exclude the RL price-based chiller demand optimisation and assume flexible demand is responding to the changes in price based on Time of Use (ToU) tariff. We consider the following two scenarios:

- Newcastle site total load with multiplying factor 50 is attached to the node 690, and
- Newcastle site total load with multiplying factor 50 is attached to the nodes 690, 670, 45, and 415 in Sydney/Newcastle area. Nodes 670, 45 and 415 have been randomly selected in the vicinity of node 690. The total load also has a constant 0.5 MW increase in demand between 5-7pm. This could represent electric vehicles or any other flexible demands that can be activated at the time of low ToU price. This means an overall increase of 25 MW aggregated demand (equivalent to 2500 EVs; 10 kW each) in each of the above 4 nodes in the sNEM 2000 model.

As it can be seen in Figure 24, and after addition of 25MW aggregated demand, the dispatch price has increased by 20% (blue spike) due to the synchronous jump in demand at 5pm (time_step=102). Note this increase has happened immediately after a sudden drop in price in the previous time step. This phenomenon is due to the same flexible demand response to the same ToU price. However, this price jump can still happen with a dynamic pricing regime, if all flexible assets respond to the dynamic price in the same synchronous way. For example, if all EV charging is based on the same smart algorithm that minimises the cost according to the same spot price predictions, the same synchronous jump in demand will happen and a jump in price is expected.



Figure 24 Impact of ToU-triggered non-scheduled flexible loads on dispatch price

Note the above dispatch price is based on the same generation cost curve for all generators regardless of their fuel types.

In the next experiment, we consider the impact of being a schedulable load on the dispatch price in the following scenarios:

- Newcastle site total load with multiplying factor 10 is attached to the node 690 as non-scheduled load, and
- Newcastle site total load with multiplying factor 10 is attached to the node 690 as scheduled load with band 1-5 as minimum demand at 0$/MW and band 6-10 as maximum demand at -10$/MW.

Note the minimum demand include all underlying loads with zero flexible demand, and the maximum demand is the total demand with full flexible load. Also note the negative bidding in band 6-10 ensures scheduling of the maximum demand if required.

As expected, and illustrated in Figure 25, the scheduled load is the same as the maximum available demand with a 10 MW peak. Therefore, the dispatch price is the same in both cases. However, if we keep increasing the multiplying factor, not all of such a cheap demand is required in the market, and hence, not all of it is dispatched. Figure 26 shows the results of the same scenario with multiplying factor 150. In this case, although the negative bidding encourages a move towards the maximum possible demand, only around 90 MW of maximum possible 150 MW is dispatched in the market. This further means the flexible cooling demand must be reduced accordingly.

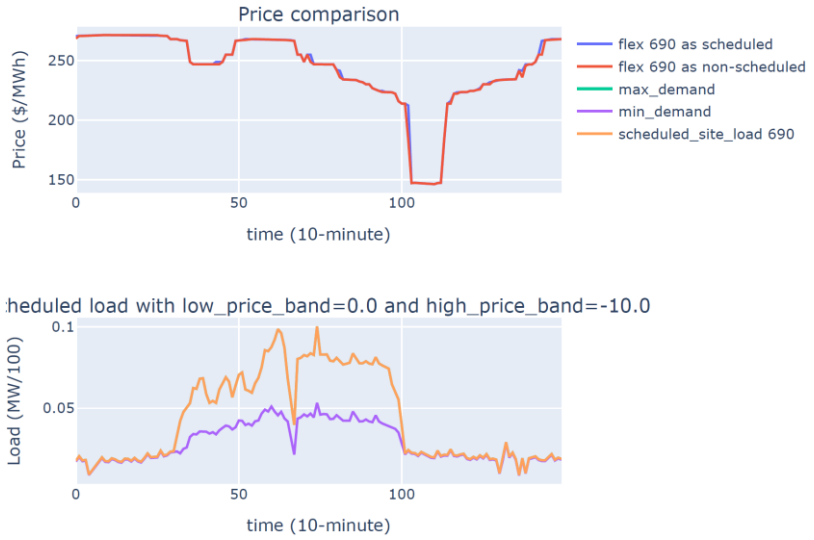


Figure 25 Impact of being schedulable load at node 690 on dispatch price with total demand multiplying factor 10



Figure 26 Impact of being schedulable load at node 690 on dispatch price with total demand multiplying factor 150

The reduction in flexible demand in Figure 26 (compared to the maximum value), causing a 5% reduction in dispatch price at around time_step=50 (minute 500 or ~8am). In other words, participation of a flexible load in the market and being schedulable can potentially reduce the spot price for everyone.

In the next experiment, we investigate the impact of scheduled load location in the following scenarios:

- Newcastle site total load with multiplying factor 150 is attached to the node 690 as non-scheduled load, and
- Newcastle site total load with multiplying factor 150 is attached to the node 605 as scheduled load with band 1-5 as minimum demand at 0$/MW and band 6-10 as maximum demand at -10$/MW.

Comparing the scheduled load prices in Figure 26 and Figure 27 reveals the importance of location of such scheduled load, where node 605 causes much lower dispatch prices than the same

scheduled load at node 690 especially in the afternoon around time_step=100 (minute 1000 or ~5pm).

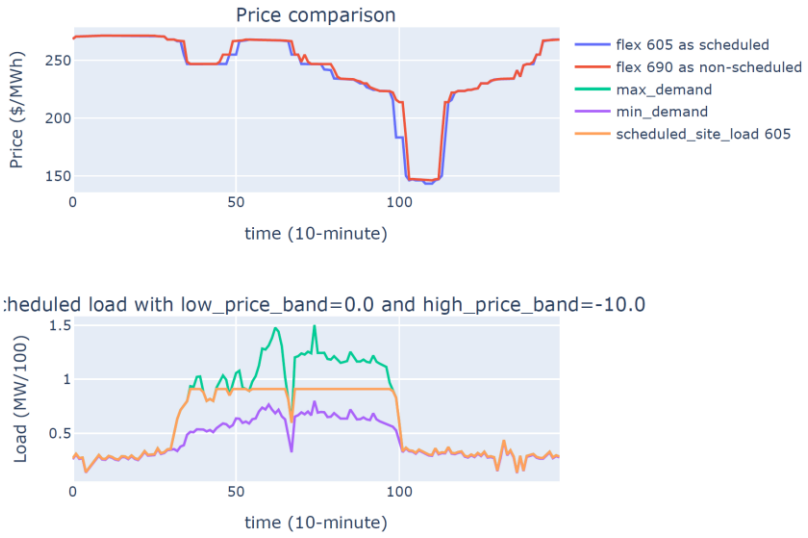


Figure 27 Impact of being schedulable load at node 605 on dispatch price with total demand multiplying factor 150

All of the above price plots in this section, so far, is based on the same generation cost curve for all generators regardless of the fuel type. If we consider the fuel-based generation bidding cost curves introduced in section 3, when scheduled load is at node 690 and multiplying factor is set at 150, we have the dispatch price in Figure 28 compared to the one in Figure 26.

Note, we still have the same conclusion with lower morning price around time_step=50 in the case of scheduled load compared to the non-scheduled load at node 690. However, there are two main differences. First, the dispatch price is completely different with lower prices during the day because of having PV generation. Second, the price signal is more volatile due to the intermittent nature of renewable generation, and third, the dispatched load is at the minimum demand towards the afternoon hours.

It is worth mentioning that the NEM market model is very complex and with a little change, for example, in the bidding curve of a scheduled load, compared to the generation bidding curves, the dispatch price may change.



Figure 28 Same scenario as in Figure 26, but with different generation cost bidding curves based on fuel types

7 Conclusions

In this project, we looked at a simplified whole system simulation from price responsive CERs located at the low-voltage (LV) level to their aggregations at the medium-voltage (MV) level, and a wholesale energy and FCAS market determining the wholesale spot price based on a synthetic National Electricity Market (NEM) model. In our proposed framework, the impact of spot price on the aggregated demand of price responsive resources (PRRs) has been studied through the identification of several price-demand Flexibility Functions (FFs).

In summary, we investigated i) the impact of variability and uncertainty of flexible loads on the dispatch price, ii) the possibility of spot price reduction by scheduling an aggregation of flexible loads, iii) the impact of size and location of flexible loads on dispatch price, iv) the impact of synchronous local optimisation (in an ensemble) on the dispatch price, and v) the impact of local optimisation goals and constraints on the dispatch price.

Future works should consider

- The use of identified price-demand flexibility function for adaptive tuning of scheduled load's price-demand bidding curve for market participation.
- The same flexibility function can also be used for the design of a desired price signal that shapes the aggregated load profiles of PRRs.
- Other types of PRRs than HVAC systems, for example, battery or EV systems.
- Updating the NEMX engine to run a unit commitment before economic dispatch.
- Updating the NEMX engine to include time-varying generation mix with time-varying generator parameters such as inertia.

References

[1] F. Arraño-Vargas and G. Konstantinou, "Modular Design and Real-Time Simulators Toward Power System Digital Twins Implementation," in IEEE Transactions on Industrial Informatics, doi: 10.1109/TII.2022.3178713 (https://doi.org/10.1109/TII.2022.3178713

[2] F. Arraño-Vargas and G. Konstantinou, "Synthetic Grid Modeling for Real-Time Simulations," 2021 IEEE PES Innovative Smart Grid Technologies - Asia (ISGT Asia), 2021, pp. 1-5, doi: 10.1109/ISGTAsia49270.2021.9715654 (https://doi.org/10.1109/ISGTAsia49270.2021.9715654

[3] R. Heidari, M. Amos and F. Geth, "An Open Optimal Power Flow Model for the Australian National Electricity Market," 2023 IEEE PES Innovative Smart Grid Technologies - Asia (ISGT Asia), Auckland, New Zealand, 2023, pp. 1-5, doi: 10.1109/ISGTAsia54891.2023.10372618.

[4] G. Mohy-ud-din, R. Heidari, H. Ergun and F. Geth, "An AC-DC Power Flow Algorithm for Australian National Electricity Market," 2023 IEEE International Conference on Energy Technologies for Future Grids (ETFG), Wollongong, Australia, 2023, pp. 1-6, doi: 10.1109/ETFG55873.2023.10407659.

[5] Ghulam Mohy-ud-din, Rahmat Heidari, Hakan Ergun, Frederik Geth,
AC–DC security-constrained optimal power flow for the Australian National Electricity Market,
Electric Power Systems Research, Volume 234, 2024, 110784, ISSN 0378-7796,
https://doi.org/10.1016/j.epsr.2024.110784.

[6] NEMX.jl, https://doi.org/10.25919/v33w-8011

[7] Young S.J., Y.V. Makarov, R. Diao, R. Fan, R. Huang, J.G. O'Brien, and M. Halappanavar, et al. 2018.**Synthetic Power Grids from Real World Models.**In *IEEE Power & Energy Society General Meeting (PESGM 2018), August 5-10, 2018, Portland, OR*. Piscataway, New Jersey:IEEE.PNNL-SA-130491.doi:10.1109/PESGM.2018.8585792

[8] Birchfield, Adam, Xu, Ti, Shetye, Komal, Overbye, ThomasACTIVSg2000: A 2000-bus synthetic grid geolocated in Texas, http://item.bettergrids.org/handle/1001/437

[9] Gillioz, M., Dubuis, G. & Jacquod, P. A large synthetic dataset for machine learning applications in power transmission grids. *Sci Data* **12**, 168 (2025). https://doi.org/10.1038/s41597-025-04479-x

[10] Wang X, Mahdavi N, Sethuvenkatraman S, West S. An environment-adaptive SAC-based HVAC control of single-zone residential and office buildings. Data-Centric Engineering. 2025;6:e3. doi:10.1017/dce.2024.57

[11] Wang, X., Mahdavi, N., Aryai, V., Sethuvenkatraman, S., and West, S. R. (October 24, 2025). "From Simulation to Reality: A Study of Reinforcement Learning Control in Operational Building Environments." ASME. J. Eng. Sustain. Bldgs. Cities. November 2025; 6(4): 041006. https://doi.org/10.1115/1.4070007

[12] Junker, R. G., Azar, A. G., Lopes, R. A., Lindberg, K. B., Reynders, G., Relan, R., & Madsen, H. (2018). Characterizing the energy flexibility of buildings and districts. *Applied Energy*, 225, 175-182.

[13] Aryai, V., Mahdavi, N., West, S. and Henze, G., 2023, November. An automated data-driven platform for buildings simulation. In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation* (pp. 61-68).

[14] Gao, G., Li, J. and Wen, Y., 2020. DeepComfort: Energy-efficient thermal comfort control in buildings via reinforcement learning. *IEEE Internet of Things Journal*, *7*(9), pp.8472-8484.

[15] Manjavacas, A., Campoy-Nieves, A., Jiménez-Raboso, J., Molina-Solana, M. and Gómez-Romero, J., 2024. An experimental evaluation of deep reinforcement learning algorithms for HVAC control. *arXiv preprint arXiv:2401.05737*.

[16] Zhang, B., Frison, L., Brox, T. and Bödecker, J., 2024. Constrained reinforcement learning for safe heat pump control. *arXiv preprint arXiv:2409.19716*.

[17] Paternain, S., Chamon, L., Calvo-Fullana, M. and Ribeiro, A., 2019. Constrained reinforcement learning has zero duality gap. *Advances in Neural Information Processing Systems*, *32*.

As Australia's national science agency, CSIRO is solving the greatest challenges through innovative science and technology.

CSIRO. Creating a better future for everyone.

Contact us

1300 363 400
+61 3 9545 2176
csiro.au/contact
csiro.au