Energy & Buildings 364 (2026) 117621

Contents lists available at ScienceDirect

Energy & Buildings

journal homepage: www.elsevier.com/locate/enbuild

Integrated renewable energy and demand-side management for low-carbon commercial buildings in tropical climates: A matched-configuration benchmark with predictive control and a multi-agent DRL architecture

Abu Hena Md Shatil [a,*], Mohammad Shakeri [b], Nowshad Amin [c], Nafiz Ahmed Chisty [a]

[a] Department of Electrical and Electronic Engineering, American International University-Bangladesh, 408/1, Kuratoli, Dhaka, 1229, Bangladesh
[b] School of Computing Technologies, RMIT University, City campus, 124 La Trobe St, Melbourne, 3000, VIC, Australia
[c] Department of Electrical Engineering, Faculty of Engineering, Islamic University of Madinah, Abo Bakr Al Siddiq Road, Madinah, 42351, Saudi Arabia

ARTICLE INFO

Keywords:
Building energy management
Model predictive control
Multi-agent reinforcement learning
Renewable energy integration
Matched-configuration baseline
Commercial buildings
Low-carbon buildings
Tropical climate

ABSTRACT

Commercial buildings account for a major share of global energy use, particularly in tropical regions where high cooling demand, intermittent renewables, and grid instability complicate effective management. A persistent confound in the renewable-enabled BEMS literature is that proposed advanced controllers are compared against rule-based controllers operating on different physical assets, conflating the contribution of the renewable hardware with that of the control strategy.

We address this by introducing a matched-configuration baseline (RBC-Full-RE) operating on the same renewable asset set as the proposed system: a 250 kWp solar PV array, 500 kW h battery, 200 m^2 solar thermal, 800 kW heat pump, and 30 kW biogas CHP. Against this matched baseline we evaluate a 12-step receding-horizon model predictive controller (MPC-Full-RE) on a 10-story, 12 500 m^2 commercial office building in Chittagong, Bangladesh, using a calibrated 3R2C thermal model and a synthesised weather year matching the local climatology.

Full-year simulation gives a clean decomposition of savings: renewable hardware contributes **28.2%** reduction in annual grid electricity (Baseline 2,030 MWh/yr → RBC-Full-RE 1,457 MWh/yr) under identical rule-based control, and MPC contributes a further **2.8 percentage points** on the same hardware (1,411 MWh/yr), totalling **30.5%** relative to the all-electric baseline. Peak demand falls 34.3% (612 kW → 402 kW); thermal comfort improves from 94.8% to 97.8% of occupied hours within the ASHRAE 55 Cat. II band. Simple payback is 12.1 years at 2024 pricing.

We additionally specify a Safe Multi-Agent DRL controller with MPC safety filtering (SMA-DRL-MPC); the MPC-Full-RE result establishes a principled lower bound on what the proposed DRL extension must improve upon. Simulation code is released for reproducibility.

1. Introduction

Commercial buildings consume approximately 40% of global primary energy and are responsible for nearly 30% of greenhouse gas emissions [1,2]. In rapidly developing tropical economies, this share is growing faster than the global average, driven by sustained electrification, persistent high temperatures and humidity that elevate cooling loads, and limited investment in efficiency-grade envelope and equipment. Bangladesh exemplifies the challenge: annual electricity demand has grown 7–10% [3], the country's coastal economic centres face frequent grid instability, and Chittagong combines high solar irradiance (5.0–5.3 kWh/m^2/day on average [4,5]) with sustained 70–85% relative humidity, producing high latent cooling loads that conventional rule-based building energy management systems (BEMS) handle poorly.

The opportunity to address this through better controls and on-site distributed energy resources (DERs) is well established in principle. Solar PV, battery storage, solar thermal, heat pumps, and biogas combined-heat-and-power (CHP) units are all individually mature technologies with established cost trajectories. The question is how to coordinate them with the building's thermal demand in a way that captures more value than rule-based scheduling. Recent advances in model predictive control (MPC) and deep reinforcement learning (DRL) for building control have shown promise in simulation, but the literature has two persistent weaknesses that motivate this work.

* Corresponding author.
E-mail address: abu.shatil@aiub.edu (A.H.M. Shatil).

https://doi.org/10.1016/j.enbuild.2026.117621
Received 6 February 2026; Received in revised form 30 April 2026; Accepted 8 May 2026
Available online 9 May 2026