What Have We Learned From Field Demonstrations of Advanced Commercial HVAC Control?

Arash J. KHABBAZI[1,2*], Elias N. PERGANTIS[1,2], Levi D. REYES PREMER[1,2], Alex H. LEE[1,2], Jie MA[1,2], Haotian LIU[1,2], Gregor P. HENZE[3,4], Kevin J. KIRCHER[1,2]

[1] School of Mechanical Engineering, Purdue University
West Lafayette, IN 47907-2099, USA

[2] Center for High Performance Buildings, Ray W. Herrick Laboratories, Purdue University
West Lafayette, IN 47907-2099, USA

[3] Department of Civil, Environmental and Architectural Engineering, University of Colorado
Boulder, CO 80309, USA

[4] National Renewable Energy Laboratory
Golden, CO 80401, USA

Contact Information: arashjkh@purdue.edu; gregor.henze@colorado.edu; kircher@purdue.edu

* Corresponding Author

ABSTRACT

Many simulation studies have suggested that advanced control strategies for heating, ventilation, and air-conditioning (HVAC) equipment in commercial buildings can help reduce energy costs and greenhouse gas emissions. However, despite these potential benefits, adoption of advanced HVAC control remains limited, due in part to a lack of confidence in the technology among decision-makers in business and government. Field demonstrations of advanced HVAC control help build confidence in the technology by demonstrating its effectiveness and economic value in the real world. This paper accordingly reviews field demonstrations of advanced commercial HVAC control strategies, such as Model Predictive Control (MPC) and Reinforcement Learning Control (RLC). This paper discusses building types, control methods, test durations, measurement and verification procedures, control objectives, and reported benefits. It further provides a critical assessment of the state of the technology and highlights research opportunities that could accelerate real-world adoption of advanced commercial HVAC control strategies. The literature review confirms that advanced HVAC controls can significantly enhance energy efficiency and occupant comfort. However, most field studies cover relatively short durations and control small spaces within larger buildings. Longer-duration studies frequently report lower savings, suggesting that short-duration studies may overestimate potential benefits. Similarly, whole-building control studies typically report lower savings than smaller-scale studies, likely because the latter tend to overlook thermal coupling between controlled zones and adjacent zones. Finally, data and discussions concerning deployment costs and challenges are almost nonexistent. This suggests an important area for future research, as achieving adoption at scale will require demonstrating not only reliable benefits but also manageable deployment costs.

1. INTRODUCTION

Residential and commercial sectors account for about 29% of total energy consumption in the United States, with nearly half of that energy is utilized to maintain thermal comfort via heating, ventilation, and air-conditioning systems (HVAC) (U.S. Energy Information Administration, 2023). Given increased environmental concerns, many HVAC manufacturers and building operators seek to enhance the efficiency of their systems. Advanced control systems can significantly improve energy efficiency and occupant comfort. They can further enable buildings to offer the flexibility of their electricity demand for power grid reliability services (Kircher et al., 2021; Priyadarshan et al., 2024). Numerous studies over the last two decades have demonstrated that advanced control algorithms, such as model predictive control (MPC), reinforcement learning control (RLC), and carefully designed rule-based control (RBC) can achieve some or all of these objectives.

Early foundational work, such as Braun (1990), focused on open-loop optimal control methods that plan a sequence of future set-points and apply them without feedback from real-time measurements or updated forecasts. Later developments by Henze et al. (1997) expanded open-loop optimal control into a comprehensive MPC framework. In MPC, models of building equipment, the envelope, and disturbances like occupancy and future weather are used. This advanced control approach applies constraints and objectives to solve an optimal control problem over a forecast horizon. The first action, such as a set-point, is implemented in the system, and the procedure is repeated at the next time-step. MPC is an established algorithm used in a wide range of control applications, often with specific guarantees about the robustness of the control system (Drgoňa et al., 2020a; Henze, 2013). Concurrently, Henze and Schoenmann (2003) investigated the use of RLC in HVAC systems. In RLC, optimal control policies are identified in an action-reward formulation, using either on-site data or a simulation model for further training data (Nagy et al., 2023). While MPC requires a mathematical model of a building's thermal dynamics, RLC enables a model-free control policy through a reward or penalty signal (Pergantis et al., 2024b).

Since the start of the twenty-first century, there has been a significant increase in advanced control studies. However, experimental research remains relatively rare, with only about 60 studies for both residential and commercial sectors. This scarcity is primarily because field studies tend to be expensive, time-consuming, and difficult. Field demonstrations are essential for improving modeling accuracy and showing building stakeholders the benefits of this technology in terms of real energy and cost savings, deployment and implementation costs, and disturbances that are not often modeled through simulations. So far, each study has been carried out independently, and there has been no coordinated effort to evaluate all of these field studies. As a result, key findings may have been overlooked, highlighting the need for a thorough investigation to find possible patterns and trends.

Several studies have partially reviewed field demonstrations of advanced HVAC controls. Sturzenegger et al. (2016) reviewed experimental building MPC research, categorizing the studies by controlled system, actuators, total experiment time, and MPC model. Blum et al. (2022) provided a chronological overview of real-world MPC implementations in commercial buildings, distinguishing them from simulations and test facility studies, and covered modeling approaches, control strategies, and performance assessment. In the context of residential buildings, Pergantis et al. (2024b) provided a chronological summary of field demonstrations for supervisory control systems in residential HVAC equipment. The current review expands upon previous works by covering a broader range of commercial buildings and provide an in-depth breakdown of key trends. To our knowledge, this paper reviews the majority of published field demonstrations of advanced supervisory commercial HVAC control systems to date. The term 'supervisory' refers to control strategies that enhance comfort and energy efficiency through real-time, predictive adjustments to HVAC systems.

This paper explains the method used to assess and analyze field demonstrations of advanced commercial HVAC control systems, followed by an overview of the relevant publications and their contributions. It presents key results from these studies, organized chronologically and categorized by control methods, building locations, test durations, objectives, and reported savings. The paper also introduces key characteristics of advanced supervisory HVAC control and analyzes their economic benefits. It concludes with a critical review of current technology and identifies research opportunities to promote advanced commercial HVAC control technologies.

2. METHODOLOGY

A comprehensive review was conducted using search engines like Google Scholar, Scopus, and journals from Elsevier, IEEE, Taylor & Francis, among others. The search covered peer-reviewed journals, conference proceedings, and technical papers on advanced control applications in commercial building HVAC systems. Notable papers also cited several useful recent studies. Given the large volume of literature, specific search strategies were required. Keywords such as "field demonstration," "practical implementation," "experimental study," "commercial building," "model predictive control," "reinforcement learning control," and "HVAC" were iteratively used to locate relevant papers. These publications were thoroughly analyzed to help shape the format and content of this review article.

The current study primarily reviews 36 field study papers. However, some of the surveyed publications investigate multiple field tests across various building types or regions (Gayeski et al., 2012; Granderson et al., 2018; Joe & Karava, 2019; Kim & Braun, 2022; Luo et al., 2022; Naug et al., 2020; Prívara et al., 2011; Široký et al., 2011; West et al., 2014; Yang et al., 2019, 2020b, 2020a; Zhang et al., 2023), accounting for a total of 56 tests. When evaluating key findings such as the reported energy and cost savings, it is crucial to consider the total number of tests rather than solely the number of papers.

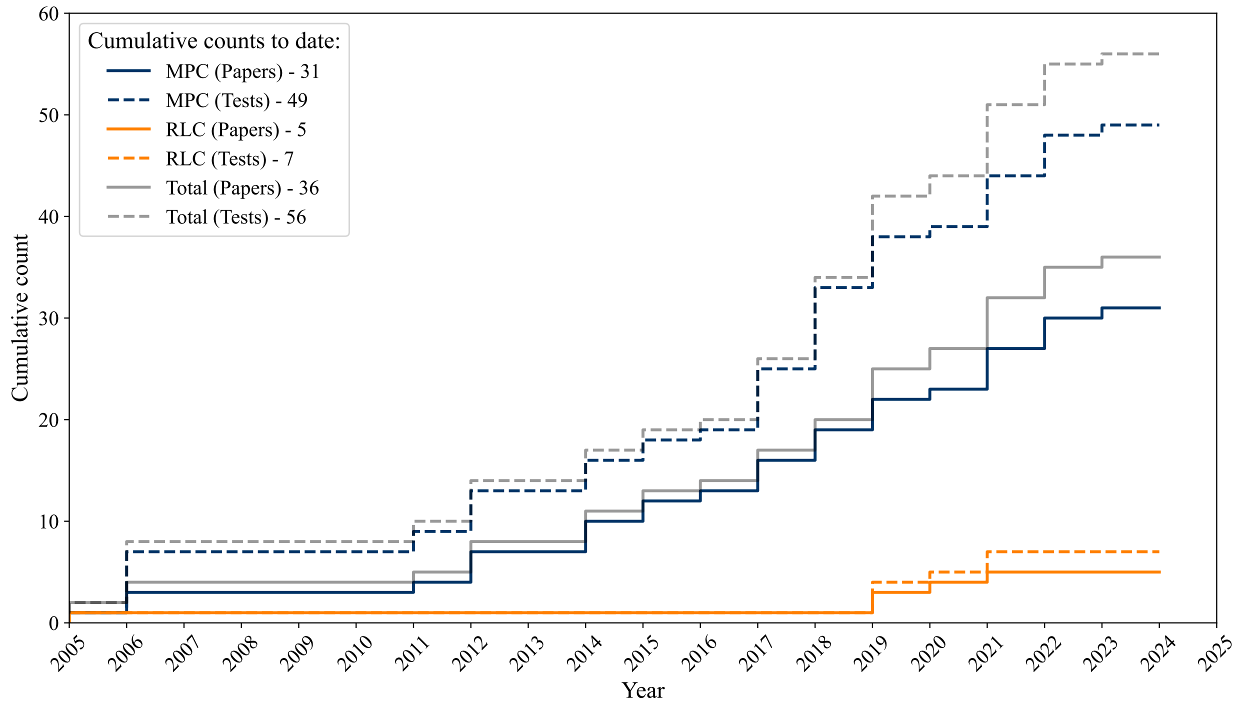

Figure 1: Yearly cumulative number of MPC, RLC, and total papers and tests. Dashed lines show the count of tests, while solid lines represent the number of publications.

Analyzing reported savings presents a challenge due to the wide range of test durations, ranging from a single day to several months. Given this, a better way is to calculate the average savings across all tests considering the data from all studies and using a duration-weighted average, as shown in Equation 1, such that

$$\overline{\text{savings}} = \frac{\sum_i (\%\ \text{savings})_i \cdot (\text{duration})_i}{\sum_i (\text{duration})_i} \tag{1}$$

where $\overline{\text{savings}}$ denotes the duration-weighted average, and the index i corresponds to each test. This strategy gives more weight to longer-duration tests, which typically yield more reliable outcomes than shorter ones.

3. RESULTS AND DISCUSSION

Figure 1 illustrates the annual cumulative number of papers published and tests conducted since 2005, categorized by MPC, RLC, and total counts. Dashed lines indicate tests, while solid lines represent publication counts. The data show a gradual increase in the number of MPC studies over the years. RLC research, deployed for the first time around 15 years ago, only experienced significant growth after 2020, which corresponds to the rapid developments in AI and reinforcement learning. This surge in RLC studies suggests a growing interest, which could be due to the growing awareness of its potential applications across various fields. This trend is expected to continue as reinforcement learning principles and applications become more widely used in the industry. For example, modern large language models like GPT-4, used by chatbots, employ reinforcement learning in their refinement stages to align responses more closely with those human critics had previously favored. Overall, the total cumulative studies show the sum of these two fields and demonstrate an upward trend, indicating rising interest in both types of control paradigms.

Multiple characteristics define a supervisory control experiment such as the control method (MPC or RLC), building location, test duration, objectives, and reported savings. To optimize space, key details are summarized chronologically and presented in Table 1 as follows:

Table 1: Summary of research on field demonstrations of supervisory commercial HVAC control.

Paper	Control method	Building type and location	Test(s) duration [days]	Objective(s)	Objective(s) improvement
Henze et al. (2005)	MPC	Test facility in Ankeny, IA, US	4	Minimize electricity and heating cost	17%–27%[1] (vs. simulation)
Liu and Henze (2006)	RLC	Test facility in Ankeny, IA, US	6	Minimize electricity and heating cost	8.3% vs. simulation
Prívara et al. (2011)	MPC	8-story university building in Prague, Czech Republic	90	Minimize energy	17%–29% vs. measurement
Široký et al. (2011)	MPC	8-story university building in Prague, Czech Republic	14–49	Minimize energy	15.5%–28.7% vs. measurement
Gayeski et al. (2012)	MPC	Test Chambers in Atlanta and Phoenix, US	7	Minimize energy	19%–25% vs. measurement
Bengea et al. (2014)	MPC	Office building in Champaign, IL, US	5	Minimize energy	60%–85% vs. measurement
Castilla et al. (2014)	MPC	Research center in Almeria, Spain	15	Minimize energy	53% vs. measurement
West et al. (2014)	MPC	3-floor office buildings in Newcastle and Melbourne, Australia	5–25	Minimize energy, cost, and GHG emissions	14%–32% energy and 16.9% cost vs. measurement
Li et al. (2015)	MPC	Office building in Philadelphia, US	20	Minimize energy	20%–33% vs. measurement
Goyal et al. (2015)	MPC	University building in Gainesville, FL, US	1	Minimize energy	40% vs. measurement
Kim et al. (2015)	MPC	Gym building affiliated with ORNL in Knoxville, TN, US	4	Minimize energy and peak demand	8% cost and 40% peak demand vs. measurement
Sturzenegger et al. (2016)	MPC	Office building in Allschwil, Switzerland	203	Minimize cost	25%[2] (vs. simulation)
De Coninck and Helsen (2016)	MPC	Office building in Brussels, Belgium	15	Minimize gas and electricity cost	20%–30% energy and 34%–40% cost vs. measurement
Hilliard et al. (2017)	MPC	University building in Halifax, Canada	113	Minimize energy	29%–63% vs. measurement
Zhuang et al. (2018)	MPC	8-floor shopping mall in Sichuan, China	1	Minimize chiller energy	16% vs. measurement
Granderson et al. (2018)	MPC	Office in Long Beach, courthouse in Dayton, hospital in Newtork, and high school in D.C., US	108–170	Minimize energy	-1.4%–8.9% vs. measurement
Kim and Braun (2018)	MPC	Retail store in Florida, US	30	Minimize cost and peak demand	12% cost and 18% peak demand vs. measurement
Joe and Karava (2019)	MPC	University building in West Lafayette, IN, US	10	Minimize energy and cost	7.8%–64% energy and 34%–78% cost vs. measurement and simulation
Lee et al. (2019)	MPC	University building in West Lafayette, IN, US	19	Minimize energy and occupant thermal dissatisfaction	occupant dissatisfaction[3] vs. measurement

Continued on next page

[1]This study compared savings in simulation (17% vs. base case and 27% vs. reference case), rather than experiment.
[2]This study compared savings in simulation (i.e. simulation to simulation), rather than using the performance data collected during testing.
[3]This study compares a self-tuned MPC controller with a baseline, showing it decreases occupant dissatisfaction by increasing energy usage.

Table 1 – *Continued from previous page*

Paper	Control method	Building type and location	Test(s) duration [days]	Objective(s)	Objective(s) improvement
Yang et al. (2019)	MPC	Office space in Singapore	2	Minimize energy	14.7%–20% vs. measurement
Yang et al. (2020b)	MPC	University building in Singapore	14	Minimize energy	16%–18% vs. measurement
Yang et al. (2020a)	MPC	University office building and lecture theatre in Singapore	10	Minimize energy	36.7%–58.5% vs. measurement
Drgoňa et al. (2020b)	MPC	Office building in Hasselt, Belgium	35	Minimize energy	53.5% vs. measurement
Naug et al. (2020)	RLC	3-floor university building in Nashville, TN, US	217	Minimize energy	7.2%–12.6% vs. measurement
Chen et al. (2020)	RLC	University conference room in Pittsburgh, US	21	Minimize energy	16.7% vs. measurement
Freund and Schmitz (2021)	MPC	Office in Hamburg, Germany	90	Minimize energy	30% vs. measurement
Touzani et al. (2021)	RLC	Test facility affiliated with LBNL in Berkeley, CA, US	7	Minimize cost	39.6% vs. measurement
Luo et al. (2022)	RLC	Several buildings in a university campus and a multi-purpose commercial building in US	90	Minimize energy	9%–13% vs. measurement
Kim and Braun (2022)	MPC	University building in West Lafayette, IN, US	5–15	Minimize cost under TOU scenarios	3.61%–8.67% vs. measurement
Blum et al. (2022)	MPC	Office affiliated with LBNL in Berkeley, CA, US	31	Minimize energy	40% vs. measurement
Kim et al. (2022)	MPC	University chiller plant in Merced, CA, US	7	Minimize peak demand and GHG emissions; improve PV-generated energy use	26.4% PV-generated energy use, 9.6% GHG emissions, and 9.8% peak demand vs. measurement
Zhang et al. (2022)	MPC	Gas station store in Blue Lake Rancheria, CA, US	N/A	Minimize cost and peak demand	11.7% cost and 34% peak demand vs. measurement
Zhan et al. (2023)	MPC	Multi-story office building in Singapore	49	Improve PV-generated energy use	19.5% PV-generated energy use and 10.6% in building self-sufficiency vs. measurement
Zhang et al. (2023)	MPC	University building in West Lafayette, IN, US	1	Minimize HVAC energy	28%–35% vs. measurement
Ham et al. (2023)	MPC	K-12 school in CA, US	30	Minimize energy and peak demand	24% peak demand and up to 16% shifted cooling/heating loads vs. measurement
Ham et al. (2024)	MPC	Small commercial building in New York, US	13	Minimize cost and peak demand	27% cost and 23% thermal load shift from peak time to off-peak time vs. measurement

The economic benefits analyzed in this study–energy reduction, cost savings, or other areas–are determined by the study's goals, current research focus, and industry demands. Current literature primarily points out energy and cost savings; thus, these form the primary focus of this analysis. Still, it is important to recognize papers that focus on peak demand reduction, such as those by Kim et al. (2015), Kim and Braun (2018), Kim et al. (2022), Zhang et al. (2022), Ham et al. (2023), and Ham et al. (2024), and those reporting the GHG emission reductions (Kim et al., 2022; West et al., 2014). Of the evaluated experiments, 41 targeted energy savings, while only 14 focused on cost savings.

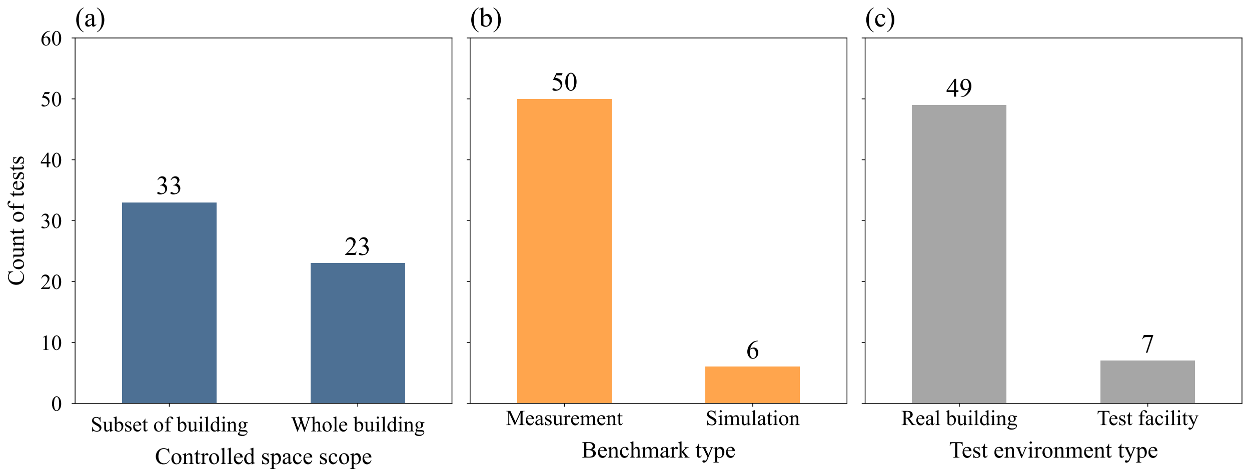

Figure 2: Bar graph comparing the number of tests carried out considering (a) controlled space scope, (b) benchmark type, and (c) test environment type.

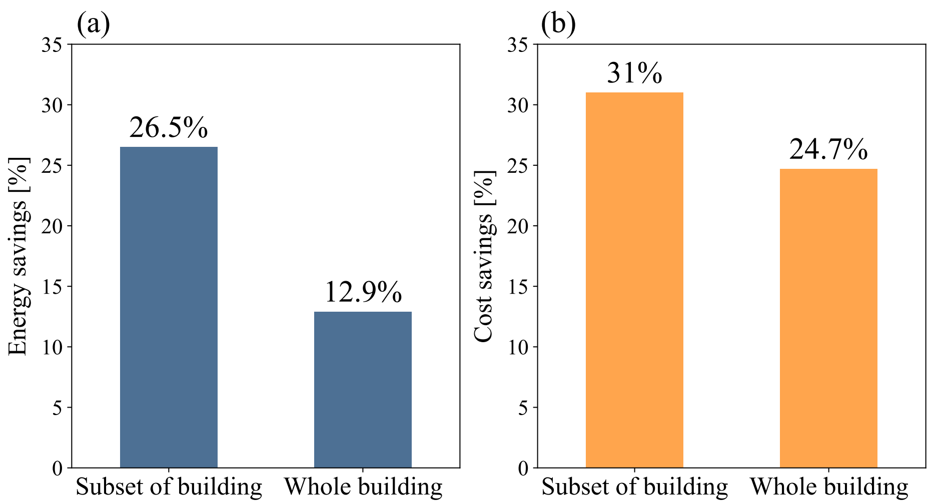

Figure 3: Duration-weighted average (a) energy and (b) cost savings, for subset of building vs. whole building. Blue bars represent the energy savings and the orange ones correspond to the cost savings.

To effectively calculate duration-weighted average savings, it is critical to to identify the key features of the experiment that most significantly influence the outcomes. Numerous factors impact this decision, including the testing environment (either real-world buildings or test facilities), the types of benchmarks used (measurement-based or simulation-based), the scope of the controlled space (whole building or a subset of the building), control modeling, controller types, objectives of the control system, HVAC delivery systems, and equipment types. Prioritizing the benchmark type, the controlled space scope, and the testing environment provides a more targeted approach to categorizing tests for calculating duration-weighted average savings, as shown in Figure 2 through three sets of bar graphs.

Specifically, Figure 2(a) shows a bar graph with both bars colored blue, comparing the number of tests using whole building control (23 tests) versus those controlling a subset of the building (33 tests). Measurement and verification (M&V) are critical in field demonstration studies since they serve as a practical benchmark for evaluating the performance of the experiments. The review of the relevant literature shows only 6 of the 56 tests used simulations, whereas the majority, 50 tests, utilized measurements to evaluate controller performance, as shown in Figure 2(b) where both bars are colored orange. This points to a clear preference for conducting field studies over simulations to validate controller developments in commercial settings.

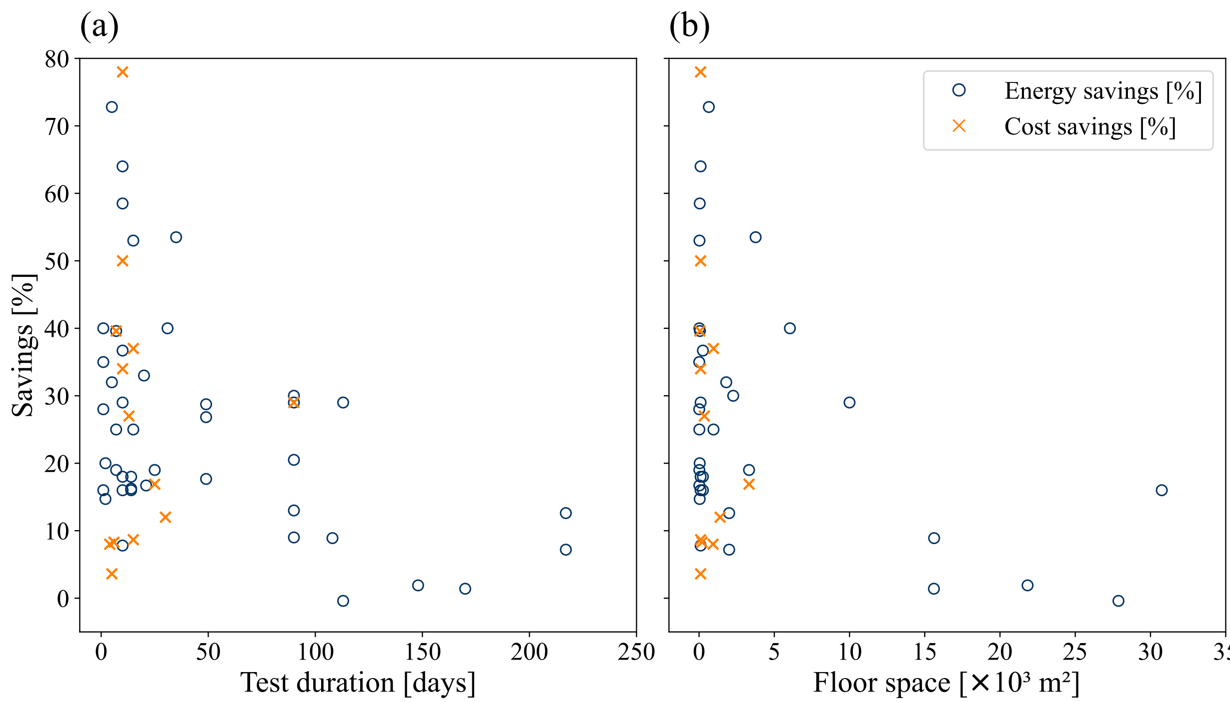

Figure 4: Duration-weighted average energy and cost savings vs. (a) test duration and (b) floor space. Blue empty circle markers show energy savings and orange cross markers indicate cost savings.

Furthermore, as depicted in Figure 2(c) with both bars colored gray, most research on supervisory commercial HVAC control systems is conducted in real-world settings such as workplaces and schools, involving actual occupants with their daily routines, with 49 tests conducted in such settings. Emphasizing these real-world scenarios is essential when assessing system effectiveness. However, the significance of test facilities in controlled experimentation and early system evaluations cannot be disregarded, as evidenced by Henze et al. (2005), Liu and Henze (2006), Gayeski et al. (2012), Yang et al. (2019), and Touzani et al. (2021), totaling 7 tests. Moving forward, including lessons learned from both real-world settings and test facilities can enhance the reliability and efficiency of advanced HVAC systems, while also addressing challenges such as occupant behavior monitoring. Given these, the controlled space's scope becomes the determining factor in calculating duration-weighted average savings, as the studies are almost evenly divided across both categories.

Figure 3 presents duration-weighted average energy and cost savings by controlled space scope. It is evident that employing supervisory HVAC control strategies has yielded substantial savings. The figure compares energy and cost savings for control strategies applied to either a subset of a building or the whole building. Blue bars indicate energy savings, with 26.5% for a subset of a building and 12.9% for the whole building, as shown in Figure 3(a). Similarly, orange bars show cost savings of 31% for a subset of a building and 24.7% for the whole building, as represented in Figure 3(b). This analysis suggests that focusing control measures on a particular subset of a building might be more effective for saving energy and costs. However, the higher savings reported for a subset of a building may not take into consideration thermal interactions between zones, thereby overestimating the savings. Future studies should look into how thermal couplings between zones influence reported savings.

Furthermore, test duration in days and floor space in m^2 were chosen to analyze the reported energy and cost savings considering their significance. Figure 4(a) illustrates a scatter plot that shows the relationship between test duration and savings. Blue empty circle markers represent energy savings, while orange cross markers indicate cost savings. Initial observations of several data points suggest an inconsistent distribution, particularly for shorter durations (less than 50 days), where savings can reach up to 80%. Although an extensive statistical analysis is required for reliable interpretations, a noticeable pattern suggests that longer test durations typically correlate with lower savings. This trend may indicate that longer testing may lead to lower energy and cost savings.

Figure 4(b) displays the relationships between energy savings, cost savings, and the floor space of the modeled zone. Most data points are concentrated in areas smaller than 5,000 m^2, showing fewer tests in larger spaces. A pattern similar to that in Figure 4(a) is observable, in which an increase in floor area is often correlated with reduced savings. This supports the findings in Figure 3, which correlate greater floor area with a shift from a subset of a building to the whole building. However, in order to perform a more comprehensive and reliable study, more features must be evaluated.

Few field demonstration studies have adequately addressed the labor, costs, and deployment challenges. Given the limited number of limited experiments to date, this oversight is not surprising. Initial pioneering efforts faced significant challenges due to the novelty of the research, limited Internet of Things (IoT) infrastructure, and the need for many forecasters to create their own weather models. Consequently, the early focus was more on demonstrating predictive smart control benefits rather than technology scalability. However, since Sturzenegger et al. (2016), studies such as Blum et al. (2022), Granderson et al. (2018), Ham et al. (2023, 2024), Kim and Braun (2018), Pergantis et al. (2024b), and Zhang et al. (2022) have systematically documented sensor costs, development costs, and deployment times. More recent papers increasingly address control solution scalability (Pergantis et al., 2024a, 2024b; Premer et al., 2024). Additionally, advancements in M&V techniques (Blum et al., 2022; Pergantis et al., 2024b) have improved the normalization of Heating Degree Days (HDD) and Cooling Degree Days (CDD), aiding in accurately assessing system performance and engaging key stakeholders, though there is still room for improvement in this area.

4. CONCLUSIONS

This paper reviewed field demonstrations of advanced supervisory control in commercial HVAC equipment, presenting five main findings. First, control technologies such as MPC and RLC can significantly enhance energy efficiency, reduce operational costs, and maintain or improve occupant comfort. Second, despite numerous simulation studies, there are comparatively few real-world field studies. This gap between simulation and field studies suggests that refocusing research efforts on field implementation could increase stakeholder confidence in the benefits and reliability of advanced control technology. Third, this paper showed (a) that experiments over shorter durations and smaller floor areas are more common, and (b) that longer and larger tests (viewed as more reliable) usually report lower savings in energy, operating costs, and greenhouse gas emissions. Fourth, this paper observed notable differences in duration-weighted average savings between tests that control a whole building and tests that control a subset of a building. Whole-building control studies often show lower savings, likely because smaller-scale studies often overlook thermal coupling between controlled zones and adjacent zones, resulting in higher perceived savings. Fifth, while every paper reviewed here discusses the benefits of advanced HVAC control, almost no papers discuss deployment labor, costs, or challenges. Future field studies should aim to fill this significant gap.

NOMENCLATURE

GHG	greenhouse gas	RLC	reinforcement learning control
HVAC	heating, ventilation, and air conditioning	RBC	rule-based control
MPC	model predictive control		

ACKNOWLEDGMENT

The Center for High Performance Buildings (CHPB) at Purdue University supported this work (project CHPB-26-2023).

REFERENCES

Bengea, S. C., Kelman, A. D., Borrelli, F., Taylor, R., & Narayanan, S. (2014). Implementation of model predictive control for an HVAC system in a mid-size commercial building. *HVAC&R Research*, *20*(1), 121–135.

Blum, D., Wang, Z., Weyandt, C., Kim, D., Wetter, M., Hong, T., & Piette, M. A. (2022). Field demonstration and implementation analysis of model predictive control in an office HVAC system. *Appl. Energy*, *318*, 119104.

Braun, J. E. (1990). *Reducing energy costs and peak electrical demand through optimal control of building thermal storage* (Anon, Ed.). Publ by ASHRAE.

Castilla, M., Álvarez, J. D., Normey-Rico, J. E., & Rodríguez, F. (2014). Thermal comfort control using a non-linear MPC strategy: A real case of study in a bioclimatic building. *J. Process Control*, *24*(6), 703–713.

Chen, B., Cai, Z., & Bergés, M. (2020). Gnu-RL: A practical and scalable reinforcement learning solution for building HVAC control using a differentiable MPC policy. *Frontiers in Built Environment*, *6*, 562239.

De Coninck, R., & Helsen, L. (2016). Practical implementation and evaluation of model predictive control for an office building in brussels. *Energy Build.*, *111*, 290–298.

Drgoňa, J., Arroyo, J., Cupeiro Figueroa, I., Blum, D., Arendt, K., Kim, D., Ollé, E. P., Oravec, J., Wetter, M., Vrabie, D. L., & Helsen, L. (2020a). All you need to know about model predictive control for buildings. *Annu. Rev. Control*, *50*, 190–232.

Drgoňa, J., Picard, D., & Helsen, L. (2020b). Cloud-based implementation of white-box model predictive control for a GEOTABS office building: A field test demonstration. *J. Process Control*, *88*, 63–77.

Freund, S., & Schmitz, G. (2021). Implementation of model predictive control in a large-sized, low-energy office building. *Build. Environ.*, *197*, 107830.

Gayeski, N. T., Armstrong, P. R., & Norford, L. K. (2012). Predictive pre-cooling of thermo-active building systems with low-lift chillers. *HVAC&R Research*, *18*(5), 858–873.

Goyal, S., Barooah, P., & Middelkoop, T. (2015). Experimental study of occupancy-based control of HVAC zones. *Appl. Energy*, *140*, 75–84.

Granderson, J., Lin, G., Singa, R., Fernandes, S., & Touzani, S. (2018). Field evaluation of performance of HVAC optimization system in commercial buildings. *Energy Build.*, *173*, 577–586.

Ham, S. W., Kim, D., Barham, T., & Ramseyer, K. (2023). The first field application of a low-cost MPC for grid-interactive K-12 schools: Lessons-learned and savings assessment. *Energy Build.*, *296*, 113351.

Ham, S. W., Paul, L., Kim, D., Pritoni, M., Brown, R., & Feng, J. (2024). Decarbonization of heat pump dual fuel systems using a practical model predictive control: Field demonstration in a small commercial building. *Appl. Energy*, *361*, 122935.

Henze, G. P. (2013). Model predictive control for buildings: A quantum leap? *Journal of Building Performance Simulation*, *6*(3), 157–158.

Henze, G. P., Dodier, R. H., & Krarti, M. (1997). Development of a predictive optimal controller for thermal energy storage systems. *HVAC&R Research*, *3*(3), 233–264.

Henze, G. P., Kalz, D. E., Liu, S., & Felsmann, C. (2005). Experimental analysis of Model-Based predictive optimal control for active and passive building thermal storage inventory. *HVAC&R Research*, *11*(2), 189–213.

Henze, G. P., & Schoenmann, J. (2003). Evaluation of reinforcement learning control for thermal energy storage systems. *HVAC&R Research*, *9*(3), 259–275.

Hilliard, T., Swan, L., & Qin, Z. (2017). Experimental implementation of whole building MPC with zone based thermal comfort adjustments. *Build. Environ.*, *125*, 326–338.

Joe, J., & Karava, P. (2019). A model predictive control strategy to optimize the performance of radiant floor heating and cooling systems in office buildings. *Appl. Energy*, *245*, 65–77.

Kim, D., Braun, J. E., Cai, J., & Fugate, D. L. (2015). Development and experimental demonstration of a plug-and-play multiple RTU coordination control algorithm for small/medium commercial buildings. *Energy Build.*, *107*, 279–293.

Kim, D., & Braun, J. E. (2018). Development, implementation and performance of a model predictive controller for packaged air conditioners in small and medium-sized commercial building applications. *Energy Build.*, *178*, 49–60.

Kim, D., & Braun, J. E. (2022). MPC solution for optimal load shifting for buildings with ON/OFF staged packaged units: Experimental demonstration, and lessons learned. *Energy Build.*, *266*, 112118.

Kim, D., Wang, Z., Brugger, J., Blum, D., Wetter, M., Hong, T., & Piette, M. A. (2022). Site demonstration and performance evaluation of MPC for a large chiller plant with TES for renewable energy integration and grid decarbonization. *Appl. Energy*, *321*, 119343.

Kircher, K. J., Cai, Y., Norford, L. K., & Leeb, S. B. (2021). Controlling big, diverse, nonlinear load aggregations for grid services by adjusting device setpoints. *2021 60th IEEE Conference on Decision and Control (CDC)*, 6377–6384.

Lee, S., Joe, J., Karava, P., Bilionis, I., & Tzempelikos, A. (2019). Implementation of a self-tuned HVAC controller to satisfy occupant thermal preferences and optimize energy use. *Energy Build.*, *194*, 301–316.

Li, P., Vrabie, D., Li, D., Bengea, S. C., Mijanovic, S., & O'Neill, Z. D. (2015). Simulation and experimental demonstration of model predictive control in a building HVAC system. *Science and Technology for the Built Environment*, *21*(6), 721–733.

Liu, S., & Henze, G. P. (2006). Experimental analysis of simulated reinforcement learning control for active and passive building thermal storage inventory: Part 2: Results and analysis. *Energy Build.*, *38*(2), 148–161.

Luo, J., Paduraru, C., Voicu, O., Chervonyi, Y., Munns, S., Li, J., Qian, C., Dutta, P., Davis, J. Q., Wu, N., Yang, X., Chang, C.-M., Li, T., Rose, R., Fan, M., Nakhost, H., Liu, T., Kirkman, B., Altamura, F., … Mankowitz, D. J. (2022). Controlling commercial cooling systems using reinforcement learning, 2022–2034.

Nagy, Z., Henze, G., Dey, S., Arroyo, J., Helsen, L., Zhang, X., Chen, B., Amasyali, K., Kurte, K., Zamzam, A., Zandi, H., Drgoňa, J., Quintana, M., McCullogh, S., Park, J. Y., Li, H., Hong, T., Brandi, S., Pinto, G., … Bernstein, A. (2023). Ten questions concerning reinforcement learning for building energy management. *Build. Environ.*, *241*, 110435.

Naug, A., Quinones-Grueiro, M., & Biswas, G. (2020). A relearning approach to reinforcement learning for control of smart buildings. *Proceedings of the Annual Conference of the Prognostics and Health Management Society, PHM*, *12*(1).

Pergantis, E. N., Premer, L. D. R., Lee, A. H., Priyadarshan, Liu, H., Groll, E., Ziviani, D., & Kircher, K. J. (2024a). Protecting residential electrical infrastructure through advanced control: The first field results. *8th International High Performance Buildings Conference at Purdue*.

Pergantis, E. N., Priyadarshan, Theeb, N. A., Dhillon, P., Ore, J. P., Ziviani, D., Groll, E. A., & Kircher, K. J. (2024b). Field demonstration of predictive heating control for an all-electric house in a cold climate. *Appl. Energy*, *360*, 122820.

Premer, L. R., Semmelmann, L., Pergantis, E. N., Groll, E. A., Ziviani, D., & Kircher, K. J. (2024). Field demonstration of predictive heat pump water heater control. *8th International High Performance Buildings Conference at Purdue*.

Prívara, S., Široký, J., Ferkl, L., & Cigler, J. (2011). Model predictive control of a building heating system: The first experience. *Energy Build.*, *43*(2), 564–572.

Priyadarshan, Pergantis, E. N., Crozier, C., Baker, K., & Kircher, K. J. (2024). Edgie: A simulation test-bed for investigating the impacts of building and vehicle electrification on distribution grids. *Proceedings of the Hawaii International Conference on System Sciences*.

Široký, J., Oldewurtel, F., Cigler, J., & Prívara, S. (2011). Experimental analysis of model predictive control for an energy efficient building heating system. *Appl. Energy*, *88*(9), 3079–3087.

Sturzenegger, D., Gyalistras, D., Morari, M., & Smith, R. S. (2016). Model predictive climate control of a swiss office building: Implementation, results, and Cost-Benefit analysis. *IEEE Trans. Control Syst. Technol.*, *24*(1), 1–12.

Touzani, S., Prakash, A. K., Wang, Z., Agarwal, S., Pritoni, M., Kiran, M., Brown, R., & Granderson, J. (2021). Controlling distributed energy resources via deep reinforcement learning for load flexibility and energy efficiency. *Appl. Energy*, *304*, 117733.

U.S. Energy Information Administration. (2023, September). Monthly energy review [Tables 1.3, 1.4c, 2.1a-2.6].

West, S. R., Ward, J. K., & Wall, J. (2014). Trial results from a model predictive control and optimisation system for commercial building HVAC. *Energy Build.*, *72*, 271–279.

Yang, S., Wan, M. P., Chen, W., Ng, B. F., & Dubey, S. (2020a). Model predictive control with adaptive machine-learning-based model for building energy efficiency and comfort optimization. *Appl. Energy*, *271*, 115147.

Yang, S., Wan, M. P., Ng, B. F., Dubey, S., Henze, G. P., Chen, W., & Baskaran, K. (2020b). Experimental study of model predictive control for an air-conditioning system with dedicated outdoor air system. *Appl. Energy*, *257*, 113920.

Yang, S., Wan, M. P., Ng, B. F., Dubey, S., Henze, G. P., Rai, S. K., & Baskaran, K. (2019). Experimental study of a model predictive control system for active chilled beam (ACB) air-conditioning system. *Energy Build.*, *203*, 109451.

Zhan, S., Dong, B., & Chong, A. (2023). Improving energy flexibility and PV self-consumption for a tropical net zero energy office building. *Energy Build.*, *278*, 112606.

Zhang, H., Tzempelikos, A., Liu, X., Lee, S., Cappelletti, F., & Gasparella, A. (2023). The impact of personal preference-based thermal control on energy use and thermal comfort: Field implementation. *Energy Build.*, *284*, 112848.

Zhang, K., Prakash, A., Paul, L., Blum, D., Alstone, P., Zoellick, J., Brown, R., & Pritoni, M. (2022). Model predictive control for demand flexibility: Real-world operation of a commercial building with photovoltaic and battery systems. *Advances in Applied Energy*, *7*, 100099.

Zhuang, J., Chen, Y., & Chen, X. (2018). A new simplified modeling method for model predictive control in a medium-sized commercial building: A case study. *Build. Environ.*, *127*, 1–12.