

Field demonstration of predictive heating control for an all-electric house in a cold climate

This is a preprint; full paper to appear in Applied Energy

Elias N. Pergantis^{a,*}, Priyadarshan^a, Nadah Al Theeb^a, Parveen Dhillon^a, Jonathan P. Ore^b, Davide Ziviani^a, Eckhard A. Groll^a, Kevin J. Kircher^{a,*}

^aCenter for High Performance Buildings, Purdue University, 177 S Russell St, West Lafayette, IN 47907, USA

^bEmerson Automation Solution, 200 Beta Drive, Pittsburgh, PA 15238, USA

Abstract

Efficient electric heat pumps that replace fossil-fueled heating systems could significantly reduce greenhouse gas emissions. However, electric heat pumps can sharply increase electricity demand, causing high utility bills and stressing the power grid. Residential neighborhoods could see particularly high electricity demand during cold weather, when heat demand rises and heat pump efficiencies fall. This paper presents the development and field demonstration of a predictive control system for an air-to-air heat pump with backup electric resistance heat. The control system adjusts indoor temperature set-points based on weather forecasts, occupancy conditions, and data-driven models of the building and heating equipment. Field tests from January to March of 2023 in an occupied, all-electric, 208 m² detached single-family house in Indiana, USA, included outdoor temperatures as low as -15 °C. On average over these tests, the control system reduced daily heating energy use by 19% (95% confidence interval: 13–24%), energy used for backup heat by 38%, and the frequency of using the highest stage (19 kW) of backup heat by 83%. Concurrent surveys of residents showed that the control system maintained satisfactory thermal comfort. The control system could reduce the house’s total annual heating costs by about \$300 (95% confidence interval: 23–34%). These real-world results could strengthen the case for deploying predictive home heating control, bringing the technology one step closer to reducing emissions, utility bills, and power grid impacts at scale.

Keywords: heat pumps, peak demand, resistance backup heat, supervisory control, predictive control

1. Introduction

1.1. Heat pumps and the power grid

Burning fossil fuels to heat buildings causes about five percent of global greenhouse gas emissions [1]. Replacing fossil-fueled heating systems by heat pumps that run on low-carbon electricity could deeply reduce these emissions. However, heat pumps can sharply increase electricity demand during cold weather, when heat demand rises and heat pump efficiencies may fall. Peaks in electricity demand can stress power grid infrastructure, increasing the risk of power outages [2]. Expanding power grid capacity can mitigate the risk of power outages,

but also incurs high costs [3], which utilities typically pass on to ratepayers [4]. For these reasons, keeping electricity reliable and affordable while accommodating rapid heat pump adoption will likely require demand management strategies as an alternative or supplement to expanding power grid capacity [5]. This paper investigates one such strategy: equipping heat pumps with supervisory control systems that can reshape electricity demand while keeping building occupants comfortable. Supervisory heat pump control systems have other potential benefits, such as improving energy efficiency, reducing utility bills, reducing emissions of greenhouse gases and other pollutants, and providing power grid reliability services [6].

*Corresponding authors: epergant@purdue.edu, kircher@purdue.edu

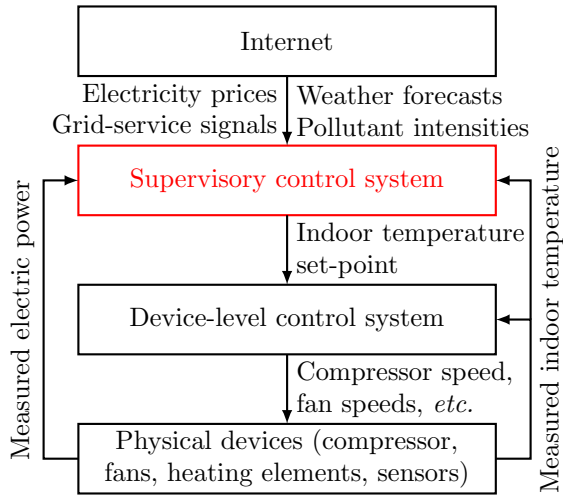


Figure 1: A heat pump control architecture. This paper develops the supervisory control system (red block).

1.2. Heat pump control systems

This paper considers the heat pump control architecture shown in Fig. 1. In this architecture, the device-level control system adjusts operating parameters, such as compressor and fan speeds, to keep the measured indoor temperature near a set-point. Device-level control systems typically use Proportional-Integral-Derivative loops with additional logic that protects equipment. This paper develops a supervisory control system that adjusts indoor temperature set-points. While sophisticated device-level control systems optimize thermal comfort and energy efficiency, they usually do not consider peak demand, dynamic prices, emissions, or grid-service revenues. Supervisory control systems can work in concert with device-level control systems to balance these higher-level objectives.

1.3. Supervisory HVAC control methods

Researchers have evaluated a variety of supervisory control methods for electric heating, ventilation and air conditioning (HVAC) equipment. Review papers such as [7, 8, 9, 10] survey these research efforts, which date back at least to the 1980s [11, 12, 13]. Seminal early contributions such as [14] focused on open-loop optimal control methods, which plan a trajectory of future set-points and implement them without feedback from real-time measurements or updated forecasts. Later, [15] and subsequent work extended the open-loop optimal control approach to a framework known as

model predictive control (MPC) [16]. Like open-loop optimal control, MPC plans a trajectory of future set-points. Unlike open-loop optimal control, MPC periodically re-plans as new measurements and forecasts arrive. In parallel with MPC research, a line of work tracing back to [17] investigated reinforcement learning control (RLC) for HVAC equipment. While MPC requires a mathematical model of a building’s thermal dynamics, RLC can learn a model-free control policy directly from a reward or penalty signal.

1.4. Simulations vs. field demonstrations

Research on supervisory HVAC control can be categorized into commercial vs. residential applications, and into computer simulations vs. field demonstrations. To the authors’ knowledge, research within the scope of this paper – residential field demonstrations – consists only of the 12 papers summarized in Table 1. For comparison, [18] recently reviewed commercial MPC field demonstrations and identified 14 published examples. Similarly, [10] recently found three published examples of commercial RLC field demonstrations. Collectively, field demonstrations comprise a small fraction of research on supervisory HVAC control systems. The review papers [7, 8, 9, 10] emphasize the importance of field demonstrations as steps toward deploying supervisory control technology at scale. Field demonstrations can highlight practical implementation challenges that require further research efforts, evaluate real-world deployment costs and labor requirements, and help convince decision-makers in business and government that the technology works and is, or could soon become, economically viable.

1.5. Contributions of this paper

This paper makes the following contributions to the literature on supervisory control of residential HVAC systems.

1. This paper provides a new long-duration field demonstration, the fifth ever to control a real home’s HVAC for more than one month.
2. This paper considers an equipment configuration that is common in North America but understudied in the research literature. Specifically, this paper considers a central air-to-air heat pump (4.5 kW rated electric power capacity) with 19 kW of backup electric resistance heating elements that turn on in stages

Table 1: Summary of field demonstrations of supervisory control systems for residential HVAC equipment.

Study, Year	Control method	Building(s), Location	Equipment	Control action(s)	Duration	Objective	Objective improvement
[19], 2013	MPC	4 houses in Denmark	Ground-to-water heat pumps	Compressor on/off state	152 days	House energy cost	9% vs. measurement
[20], 2014	MPC	1 house in the USA	Resistance radiant panels, air conditioner	Indoor temperature set-point, air flow	10 days	HVAC energy	18–30% vs. simulation
[21], 2017	MPC	1 house in Canada	Ground-to-water heat pump	Indoor temperature set-point	22 days	HVAC energy cost	16–50% vs. simulation
[22], 2020	MPC	1 bedroom in Switzerland	Radiant panels	Water flow	6 days	HVAC energy	25% vs. measurement
[23], 2020	MPC	1 house in the Netherlands	Air-to-water heat pump	Indoor temperature set-point	3 days	HVAC energy cost	12–15% vs. measurement
[24], 2020	RLC	1 house in the USA	Central air conditioner	On/off state and indoor temperature set-point	5 days	HVAC energy cost	7–17% vs. simulation
[25], 2021	MPC	1 house in Norway	Radiator	Indoor temperature set-point	14 days	HVAC energy cost	22% vs. simulation
[26], 2022	MPC	2 bedrooms in Switzerland	Radiant panels	Water flow	156 days	HVAC energy	26–49% vs. measurement
[27], 2022	MPC	1 house in Italy	Air-to-water heat pump	Compressor and pump speeds	5 days	HVAC energy cost	10–17% vs. measurement
[28], 2023	MPC	2 rooms in China	Mini-split air conditioners	On-off state	12 days	HVAC energy cost	22–27% vs. measurement
[29], 2023	MPC	1 house in Denmark	Air-to-water heat pump	Valve position, outdoor temperature override	97 days	HVAC energy, emission cost	2–17% vs. measurement
[30], 2023	MPC	1 solar house in Canada	Radiant floors	Pump on/off states	182 days	Temperature regulation	71% vs. measurement

when the heat pump cannot keep up with heat demand. While three residential field demonstrations in the literature considered air-to-air cooling equipment [20, 24, 28], none studied air-to-air heating or resistance backup heat. These missing elements constitute an important research gap, since (a) installers in North America often pair heat pumps with resistance backup heat, especially in cold climates; and (b) this equipment configuration can lead to high peaks in electricity demand in cold weather, when heat demand rises, heat pump efficiencies may fall, and inefficient resistance backup heat may turn on. To handle this equipment configuration within a convex optimization framework, this paper introduces a convex reformulation of the mapping from the heating plant’s combined thermal power output to its electrical power input.

3. This paper develops and field-tests a new, adaptive method for dynamically balancing

the competing objectives of energy efficiency and thermal comfort. Using predictions of user preferences and weather conditions, this method periodically adjusts a weight on a thermal discomfort term in the predictive controller’s objective function. Field tests, which included regular surveys of real occupants’ thermal comfort, showed that this method kept occupants comfortable while yielding substantial energy cost savings.

4. This paper presents the first residential field demonstration to include reducing peaks in electricity demand as a control objective or to significantly reduce demand peaks. This field evidence builds confidence that advanced HVAC control can mitigate capacity expansion requirements for electrical infrastructure.
5. This paper presents a comprehensive literature review of all published field demonstrations of supervisory control of residential HVAC, identifying several knowledge gaps.

6. This paper discusses implementation difficulties with supervisory HVAC control, strategies for overcoming those difficulties, and ways to improve scalability by mitigating hardware and labor requirements. While many studies discuss the *benefits* of supervisory HVAC control, very few discuss the *costs*, which strongly influence economic viability. This paper provides a real-world data point on deployment cost that can inform future benefit/cost analyses.

1.6. Organization of this paper

Section 2 of this paper reviews prior field demonstrations of supervisory control for residential HVAC equipment. Section 3 discusses the building, hardware, and software used in our field demonstration. Section 4 develops the modeling, learning, and control methods underlying the supervisory control system. Section 5 presents field demonstration results. Section 6 discusses practical implementation challenges and possible ways to improve scalability. Appendix A defines acronyms and mathematical notation. Appendix B discusses other possible formulations of the optimal control problem considered in §4.

2. Review of residential field demonstrations

Table 1 summarizes 12 field demonstrations of supervisory residential HVAC control. To the best of the authors’ knowledge, Table 1 includes all residential field demonstration studies published to date. Seven of the 12 studies in Table 1 took place in Europe, two in the USA, two in Canada, and one in China. Ten of the 12 studies controlled detached single-family homes, while [26] and [28] each controlled two rooms within larger buildings. Nine of the 12 studies controlled hydronic heating systems, while [20] and [24] controlled central air conditioners and [28] controlled mini-split air conditioners. Control actions varied across studies, with some directly manipulating equipment speeds or on/off states and others manipulating set-points of temperatures or fluid flows. Experiment durations ranged from three to 182 days. Eleven of the 12 studies optimized for energy or energy cost; [29] also optimized for greenhouse gas emissions, while [30] focused on improving temperature regulation in a complex HVAC system for a solar home. Savings in energy or energy cost varied from 2–49%, with wide variation in performance evaluation practices.

This section provides more detail on methods, performance evaluation, and research gaps.

2.1. Methods

2.1.1. Thermal modeling

Of the studies listed in Table 1, all but one [24] used MPC. All studies modeled thermal dynamics using low-order linear thermal circuit models. This includes [24], which used a thermal circuit model to train an RLC policy offline. The majority of studies used models with 1–3 thermal capacitances and 1–4 resistances. These studies identified model parameters using linear regression, genetic algorithms, particle swarm optimization, or subspace methods for linear state-space models. The amount of training data varied from five days [28] to 370 days [26], with most studies using about one month of data. These trends broadly corroborate the findings in [31] that low-order thermal circuits trained on 3–5 weeks of data typically perform well enough for control.

2.1.2. Disturbance prediction

Disturbances are uncertain, exogenous inputs that influence a system’s evolution. Accurate disturbance predictions can significantly improve control performance. The studies in Table 1 obtained weather predictions from weather services or forecast models trained on on-site measurements [20, 23]. Some studies predicted solar heat gains using physics-inspired methods that consider solar angles and optical properties of materials [29, 30]. Other studies predicted solar photovoltaic electricity generation [27] or wholesale electricity market prices [23, 29]. Only one study [20] predicted occupancy.

2.1.3. Sensing

Unlike commercial buildings, which often have extensive monitoring and control systems, residential buildings typically have limited sensing infrastructure. With the exception of [28], which used Internet-connected outlets with built-in electric power sensors, the studies reviewed in Table 1 all measured the thermal power¹ supplied by HVAC equipment. Thermal power sensing typically involves both low-cost temperature sensors and more expensive air or water flow sensors. Two studies [20, 23] measured local weather conditions.

¹This paper uses ‘thermal power’ as a shorthand for the rate of heat transfer to or from a thermodynamic system.

2.1.4. Control actions

In the system architecture illustrated in Fig. 1, the control action is the signal that the supervisory control system sends to the device-level control system. In the studies in Table 1, most supervisory control systems adjusted the temperature set-point of the indoor air or the water in a storage tank. Some studies adjusted lower-level variables, such as the position of a water flow valve, the on/off state of a heat pump or air conditioner, or the speed of a compressor or pump. In one study, the supervisory control system adjusted a heat pump’s perceived outdoor temperature by manipulating the voltage signal from a sensor [29].

2.1.5. Control objectives

All but one of the studies in Table 1 aimed to minimize the energy used by HVAC equipment or the cost of HVAC energy. All studies aimed to promote thermal comfort by maintaining the indoor temperature within a given band [20, 21, 30] or by penalizing deviations of the indoor temperature from the occupants’ preferred temperature [29]. Other objectives included maximizing on-site use of solar photovoltaic electricity [27] and minimizing greenhouse gas emissions under a time-varying greenhouse gas intensity of electricity [29].

2.1.6. Thermal comfort

Most studies in Table 1 sought to maintain occupants’ thermal comfort by constraining indoor air temperatures within a band around an occupant-specified preference. No study analyzed thermal comfort using more detailed models such as Predicted Percentage Dissatisfied (PPD) [32, 33]. (PPD models the expected percentage of people who would express discomfort under a given set of conditions, such as air temperature, air flow, clothing level, and activity level.) Only five of the twelve studies controlled spaces with more than a few hours per day of occupancy [19, 22, 23, 26, 29]. Of these, only [29] reported on occupants’ comfort during testing. In [29], wider temperature constraint bands led to higher energy savings.

2.2. Performance evaluation and savings estimates

The 12 studies in Table 1 reported a wide range of energy and cost savings, from 2–50%. Some of this variation follows from differences in control algorithms, HVAC equipment, building envelopes, and climates. However, experiment setups and reporting practices also vary in several important ways.

Table 2: Duration-weighted average energy cost savings

Compared to measurements?	Controlled whole buildings?	Reported savings
Yes	Yes	15%
No	Yes	26%
Yes	No	36%

2.2.1. Long vs. short experiments

Only five of the 12 studies lasted more than two weeks [19, 21, 26, 29, 30]. Four of the 12 studies lasted less than one week. Longer experiments yield larger sample sizes and less uncertainty in savings estimates.

2.2.2. Comparison to measurement vs. simulations

Eight of the 12 studies compared the measured performance of an advanced supervisory controller to the *measured* performance of a benchmark controller [19, 22, 23, 26, 27, 28, 29, 30]. In our experience, this approach typically introduces less uncertainty than comparing to the *simulated* performance of a benchmark controller. Even the highest-performing building models often have energy prediction errors of 15% or more – model ‘noise’ of similar scale to the savings ‘signal.’ Comparing the measured performance of two controllers brings its own challenges, as weather and occupant behavior vary from day to day. Normalizing performance metrics by weather statistics, such as heating or cooling degree-days, facilitates comparisons across disparate weather conditions.

2.2.3. Whole buildings vs. rooms

Three of the 12 studies controlled individual rooms within larger buildings [22, 26, 28]. This approach introduces an accounting risk. In winter, for example, lowering the air temperature in one room will cause heat to flow into it from adjacent, warmer rooms. Maintaining comfortable temperatures within the adjacent rooms will therefore require more energy. Reporting only the energy savings in the controlled room, without accounting for increased energy used to heat adjacent rooms, overestimates the total savings.

2.2.4. Current state of reliable savings estimates

Of the 12 studies in Table 1, only [19, 29, 30] reported experiments that lasted longer than two weeks, compared controller performance to the measured performance of a benchmark controller,



Figure 2: The DC Nanogrid House is a 208 m², 1920s-era house with all-electric appliances in West Lafayette, Indiana, USA.

and controlled whole buildings. Of those, [19] reported 9% whole-house energy cost savings over 152 days, [29] reported 2–17% savings in the combined cost of electricity and greenhouse gas emissions over 97 days, and [30] did not report savings because the study aimed to improve indoor temperature regulation in a complex HVAC system. Assuming HVAC equipment used about half of whole-house energy, the reported 9% whole-house energy cost savings in [19] translate to about 18% HVAC energy cost savings. The two short-duration studies that compared to measurements and controlled whole buildings reported HVAC energy cost savings of 12–15% over three days [23] and 10–17% over five days [27]. The weighted average of the midpoints of the HVAC energy cost savings ranges reported in [19, 23, 27, 29], weighted by experiment duration, is 15%.

2.2.5. Savings overestimation tendencies

The duration-weighted average of the midpoint savings estimates in studies that controlled whole buildings, but compared controller performance to simulations rather than measurements, is 26% [20, 21, 24, 25]. Similarly, the duration-weighted average in studies that compared controller performance to measurements, but controlled individual rooms rather than whole buildings, is 36% [22, 26, 28]. Relative to the 15% duration-weighted average in studies that compared performance to measurements and controlled whole buildings, these estimates are higher by factors of 1.7 and 2.4, respectively. Table 2 summarizes these calculations. While some of the discrepancies in reported savings could trace back to differences in algorithms, prices, buildings, equipment, or climates, the discrepancies are large enough to suggest the possibility of systematic overestimation due to comparing

performance to simulations, controlling individual rooms within larger buildings, or both.

Future work could use statistical methods to formally test the hypotheses that comparing to simulations and/or controlling individual rooms within larger buildings leads to systematic overestimation of savings. If a thorough analysis supported either hypothesis, the research community could propose improved standards for measurement and verification. For example, standards could recommend comparing controller performance only to measurements, not simulations; or comparing to simulations but presenting savings estimates alongside confidence intervals that quantify the accuracy of benchmark models. Standards might also recommend either controlling whole buildings or explicitly accounting for heat transfer between controlled rooms and adjacent thermal zones in savings estimates.

2.3. Research gaps

Reviewing all available research literature on field demonstrations of supervisory control for residential HVAC equipment identified three knowledge gaps, all of which this paper helps to fill.

2.3.1. Lack of reliable savings estimates

As discussed above, only two field studies [19, 29] demonstrated a supervisory residential HVAC controller for more than two weeks, compared its performance to the measured (rather than simulated) performance of a benchmark controller, and either controlled a whole building (rather than rooms within larger buildings) or accounted for increased energy used to condition adjacent rooms. Additional studies that meet these criteria would refine estimates of HVAC energy cost savings. Refined savings estimates could help businesses decide

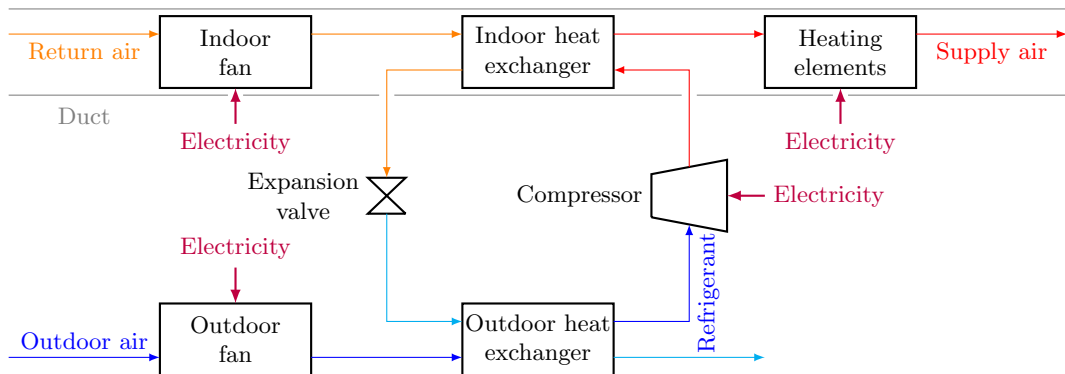


Figure 3: The DC Nanogrid House’s HVAC equipment in heating mode. The colors indicate cold, cool, warm, and hot.

whether to invest in creating supervisory HVAC control products. Refined savings estimates could also help government agencies decide whether and how to incentivize adoption of supervisory HVAC control systems.

2.3.2. Lack of deployment cost and labor estimates

Deployment costs and labor influence the scalability of supervisory control technology. The residential field demonstration studies in Table 1 contain relatively little discussion of deployment costs or labor. Aside from [28], which quoted a total hardware cost of 350 Chinese Yuan (about 50 USA dollars), no studies reported how much their sensing, communication, or actuation hardware cost. Similarly, no studies discussed how many hours of engineer labor their field deployments required, or how the labor requirements broke down between tasks. While [18] and [34] discussed deployment costs and labor for commercial buildings, Section 6 of this paper provides the first thorough discussion in the residential context.

2.3.3. Lack of diversity in objectives and equipment

Most studies in Table 1 sought to minimize energy costs in a single-device, single-source HVAC system. While energy costs are important, supervisory HVAC control systems can provide other benefits that the field demonstrations in Table 1 did not consider. For example, no studies sought to minimize peak electricity demand or maximize flexibility for power grid reliability services. Similarly, no studies considered dual-source HVAC systems, such as electric heat pumps with backup heat from natural gas [35] or electric resistance. Aside from [27], which coordinated a real heat pump with simulated solar photovoltaics, no studies coordinated

HVAC systems with other devices. This leaves a gap between field demonstrations and simulation studies that have coordinated HVAC systems with batteries and solar photovoltaics [36, 37, 38], electric vehicles [39], and other devices [40, 41]. A recent field demonstration in a commercial building showed promise in coordinating HVAC systems with other devices [42].

2.3.4. Treatment of occupant thermal comfort

Only five studies controlled spaces in residential buildings with more than a few hours per day of occupancy [19, 22, 23, 26, 29]. Of these, only [29] reported on occupants’ comfort during testing or investigated the relationship between thermal comfort and energy costs, pollutant emissions, demand peaks, or other objectives. To the authors’ knowledge, no study has yet developed methods to adaptively balance trade-offs between thermal comfort and other objectives under dynamic occupant preferences and weather conditions.

3. Building, hardware, and software

3.1. The DC Nanogrid House

The field demonstrations in this paper took place in the DC Nanogrid House, pictured in Fig. 2. The DC Nanogrid House is a 208 m², two-story, 1920s-era detached single-family home near Purdue University’s campus in West Lafayette, Indiana, USA. This location falls under the International Energy Conservation Climate Code 5A. This climate sees both hot, humid summers (up to 35 °C) and cold winters (down to −20 °C). We have renovated the DC Nanogrid House with envelope upgrades and all-electric appliances [43], including a central air-to-air heat pump with backup resistance heating

elements. The exterior walls have foam insulation with an R-Value of $3.5 \text{ }^\circ\text{C m}^2/\text{W}$. Code-minimum U-8 $\text{W}/\text{m}^2/^\circ\text{C}$ windows make up about 20% of the exterior wall area. In addition to conventional alternating current wiring, the DC Nanogrid House has a direct current (DC) electrical network that connects its rooftop solar panels, home battery, and DC appliances [44]. We have installed networked sensors throughout the thermal and electrical systems. The DC Nanogrid House is a living laboratory occupied by graduate engineering students.

3.2. HVAC equipment

An air-to-air heat pump, illustrated in Fig. 3, provides space conditioning for the DC Nanogrid House. The heat pump has 14 kW of rated cooling capacity, a cooling seasonal coefficient of performance (seasonal COP) of 5.3, and a heating seasonal COP of 2.5. Up to 19.2 kW of backup heating elements turn on in stages (9.6, 14.4, and 19.2 kW) when the heat pump’s vapor-compression cycle cannot keep up with heating demand. The heat pump’s indoor heat exchanger, the heating elements, and an indoor fan reside within a duct in the basement. The indoor fan blows conditioned air through ducts to 12 vents within the house, drawing air back from six return vents. A second fan blows air over the outdoor heat exchanger. The HVAC system has no outdoor air intake, as infiltration through gaps in the building envelope keeps indoor air fresh. Both fans and the heat pump’s compressor have variable-speed drives. To keep the measured indoor temperature near its set-point, the device-level control system (see Fig. 1) modulates both fan speeds and the compressor speed roughly in unison. The supervisory control system developed here decides indoor temperature set-points and sends them to the device-level control system.

3.3. Sensing, communication, and computing

The supervisory control system uses measurements from Internet-connected power and temperature sensors. An IoTaWatt electric power meter measures a common voltage, as well as individual currents on three circuits that serve the compressor, fans, and heating elements. The electric power meter communicates via Wi-Fi to an InfluxDB database hosted by DigitalOcean. The database also stores the heat pump thermostat’s indoor temperature measurement, extracted via Wi-Fi. A desktop computer, located at the DC

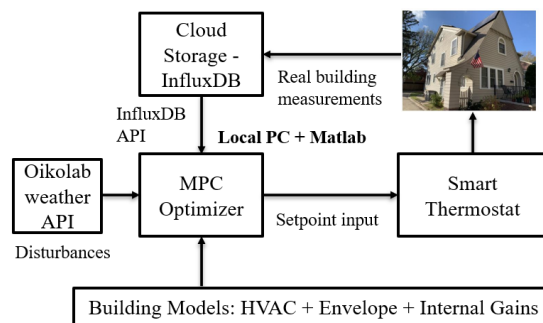


Figure 4: Information flow in the control system.

Nanogrid House, periodically pulls measurements from the database, downloads day-ahead hourly weather forecasts from Oikolab, computes the next indoor temperature set-point in Matlab using the CVX toolbox [45], and pushes the set-point to the thermostat via Wi-Fi and an application programming interface (API). Fig. 4 illustrates information flow in the cyber-physical systems.

The sensing, communication, and computing infrastructure described above suffices for supervisory control in practice. For research purposes, we also installed airflow and temperature sensors in the duct upstream of the indoor fan and air temperature sensors downstream of the heating elements (see Fig. 3). These sensors provide an estimate of the HVAC plant’s thermal power output. The airflow and temperature sensors communicate on serial wires using the Modbus protocol. The wires connect to a Raspberry Pi (a single-board computer with attached input/output circuit pins), which pushes data via Wi-Fi to the InfluxDB database. More information on the sensors and networking can be found in [46, 47].

4. Modeling, learning, and control methods

This section describes the building and HVAC equipment models, the model training methodology, and the MPC formulation. The modeling approach proposed here combines physical intuition with supervised learning techniques for time-series regression. This fusion of physics with machine learning significantly reduces data requirements relative to purely data-driven modeling approaches. In the field, this approach yielded a plant model with sufficient accuracy for control purposes after gathering one month of training data.

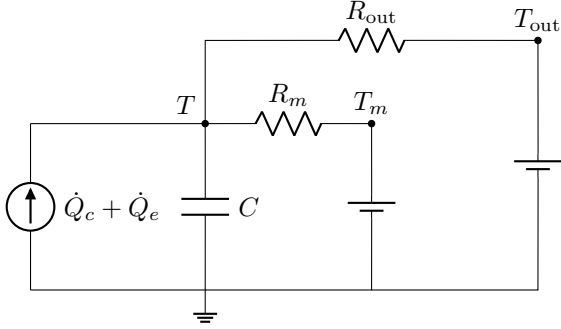


Figure 5: Thermal circuit model of the DC Nanogrid House.

4.1. Building modeling

4.1.1. Thermal circuit model

An equivalent thermal circuit model, illustrated in Fig. 5, captures the DC Nanogrid House’s temperature dynamics with sufficient accuracy for control purposes. Thermal circuits behave similarly to electrical circuits, with temperature playing the role of voltage and thermal power playing the role of current. The model in Fig. 5 has one state: the temperature T ($^{\circ}\text{C}$) of the indoor air and shallow thermal mass, such as ducts and the surfaces of walls, floors, and furniture. The model has four input signals: the temperature T_m ($^{\circ}\text{C}$) of the deep thermal mass, such as the interior portions of walls, floors, and furniture; the outdoor air temperature T_{out} ($^{\circ}\text{C}$); the thermal power \dot{Q}_c (kW) from controlled HVAC equipment; and the thermal power \dot{Q}_e (kW) from exogenous sources such as lights, appliances, electronics, bodies, and the sun. The model has three parameters: the indoor thermal capacitance C (kWh/ $^{\circ}\text{C}$), the thermal resistance R_m ($^{\circ}\text{C}/\text{kW}$) between the indoor air and deep thermal mass, and the thermal resistance R_{out} ($^{\circ}\text{C}/\text{kW}$) between the indoor and outdoor air.

4.1.2. Governing equations

The indoor temperature dynamics follow the first-order linear ordinary differential equation

$$C \frac{dT}{dt} = \frac{T_m - T}{R_m} + \frac{T_{\text{out}} - T}{R_{\text{out}}} + \dot{Q}_c + \dot{Q}_e. \quad (1)$$

Defining the effective boundary temperature

$$\theta = \frac{R_{\text{out}}T_m + R_mT_{\text{out}}}{R_m + R_{\text{out}}} \quad (2)$$

and the effective thermal resistance

$$R = \frac{R_mR_{\text{out}}}{R_m + R_{\text{out}}} \quad (3)$$

gives an equivalent differential equation:

$$C \frac{dT}{dt} = \frac{\theta - T}{R} + \dot{Q}_c + \dot{Q}_e. \quad (4)$$

Assuming the input signals θ , \dot{Q}_c , and \dot{Q}_e are constant over each time step of uniform duration Δt (h), the exact discrete-time dynamics are

$$T(k+1) = aT(k) + (1-a)[\theta(k) + R(\dot{Q}_c(k) + \dot{Q}_e(k))], \quad (5)$$

where the integer k indexes time steps and

$$a = \exp\left(-\frac{\Delta t}{RC}\right).$$

The discrete-time dynamics (5) come from analytically solving to the first-order linear ordinary differential equation (4) that governs the continuous-time dynamics.

4.2. Building model training

The supervisory control system measures the indoor temperature T , downloads measurements of the outdoor temperature T_{out} from a weather service, and either measures the HVAC thermal power \dot{Q}_c or calculates it as the product of the measured electric power and the heat pump COP. Therefore, training the thermal circuit model requires instantiating the deep thermal mass temperature T_m , the thermal resistances R_{out} and R_m , the discrete-time dynamics parameter a (or thermal capacitance C), and the exogenous thermal power \dot{Q}_e . We trained the model on passive observations from November 11 to December 10, 2022. These 30 days gave 720 hourly data points for each variable.

4.2.1. Deep mass temperature

Because deep thermal mass temperatures typically change significantly slower than indoor air temperatures, we modeled T_m as constant. This simplification reduced the number of parameters to identify and avoided the need to estimate unmeasured mass temperatures. In buildings with more thermal mass activation, for example through radiant floor heating, the control system could estimate unmeasured thermal mass temperatures using a Kalman filter. However, the constant temperature assumption led to acceptable controller performance with the forced-air heating system in the DC Nanogrid House, consistent with other studies of residential forced-air systems [28, 31]. Finally, the modeling approach developed here pre-

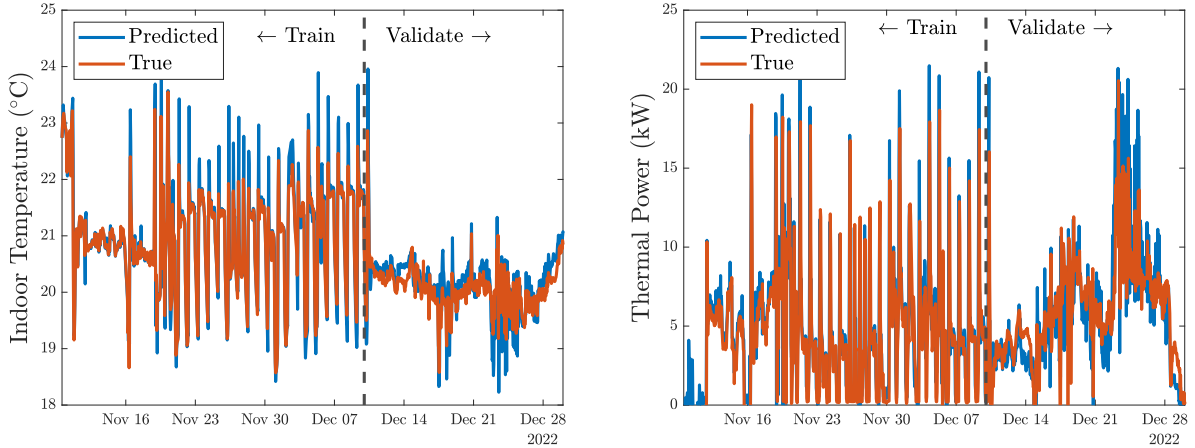


Figure 6: The full model’s predictions (blue curves) of the indoor temperature (left) and the HVAC thermal power (right) match the targets (orange curves) reasonably well in both training (left two-thirds of data) and validation (right one-third). The validation root-mean-square temperature and power prediction errors were 0.41 °C and 2.3 kW, respectively.

dicts a catch-all disturbance term that includes unmodeled dynamics such as heat transfer with time-varying thermal mass (see Section 4.2.4 for details). As most deep thermal mass in the DC Nanogrid House resides in interior material, rather than exterior walls, we estimated the thermal mass temperature by the time-average indoor temperature in the training data, 20.6 °C.

4.2.2. Outdoor resistance

In steady states with $T_m \approx T$ and approximately constant \dot{Q}_e , the dynamics reduce to

$$T(k) - T_{\text{out}}(k) \approx \alpha + \dot{Q}_c(k)R_{\text{out}}, \quad (6)$$

where $\alpha \approx R_{\text{out}}\dot{Q}_e$. To estimate R_{out} , we formed a dataset from steady times k_1, \dots, k_n during nights, when \dot{Q}_e was approximately constant. The dataset has targets $T(k_i) - T_{\text{out}}(k_i)$, features $(1, \dot{Q}_c(k_i))$, and model

$$\begin{bmatrix} T(k_1) - T_{\text{out}}(k_1) \\ \vdots \\ T(k_n) - T_{\text{out}}(k_n) \end{bmatrix} \approx \begin{bmatrix} 1 & \dot{Q}_c(k_1) \\ \vdots & \vdots \\ 1 & \dot{Q}_c(k_n) \end{bmatrix} \begin{bmatrix} \alpha \\ R_{\text{out}} \end{bmatrix}. \quad (7)$$

Fitting this model via least squares gave an R_{out} estimate of 2.04 °C/kW.

4.2.3. Capacitance and mass resistance

We fit the capacitance C and resistance R_m using linear regression with a grid search over R_m values from 0.01 to 10 °C/kW. Each value of R_m defines an effective boundary temperature θ and resistance R

via Eq. (2) and (3). Given θ and R , the dynamics during unsteady periods with approximately constant \dot{Q}_e give

$$\begin{aligned} T(k+1) - \theta(k) - R\dot{Q}_c(k) \\ \approx \beta + (T(k) - \theta(k) - R\dot{Q}_c(k))a, \end{aligned} \quad (8)$$

where $\beta \approx (1-a)R\dot{Q}_e$. As with the outdoor resistance, we fit this linear model via least squares, giving an estimate of a for each value of R_m . We split the available unsteady data into training and validation sets. We found that $R_m = 1.06$ °C/kW and the corresponding $a = 0.8$ generalized well to the validation data. Although the supervisory control system does not use the capacitance $C = -\Delta t / (R \ln(a))$, the estimate of 6.5 kWh/°C, while somewhat higher than expected, matched our physical intuition reasonably well.

4.2.4. Exogenous thermal power prediction

After fitting the model parameters, we inverted the dynamics to form a dataset with targets

$$\dot{Q}_e(k) = \frac{1}{R} \left(\frac{T(k+1) - aT(k) - \theta(k)}{1-a} \right) - \dot{Q}_c(k). \quad (9)$$

We then followed a standard supervised learning workflow for time-series prediction. We evaluated features including the outdoor temperature, global solar irradiance, wind speed, time of day, day of week, and past values of the targets. We compared model structures including linear autoregression, neural networks, regression trees, and support

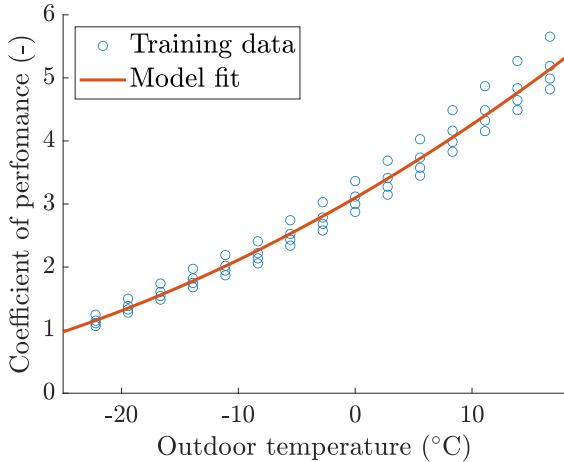


Figure 7: The COP model depends only on the outdoor temperature, but matches manufacturer data with $R^2 = 0.98$.

vector machines. We tuned model hyperparameters using two-fold cross-validation. A support vector machine with all of the weather and time features, but no autoregression, gave an acceptable balance between model complexity and prediction accuracy.

4.2.5. Trained model performance

Fig. 6 shows the end-to-end fit of the full model, including the thermal circuit and the exogenous thermal power predictor, in training and validation data. The one-step-ahead indoor temperature predictions (left plot) match the targets with a root-mean-square error of $0.41\text{ }^\circ\text{C}$ in the validation data, which we did not use for model training. The HVAC thermal power predictions, formed by substituting the indoor temperature measurements and exogenous thermal power predictions into the dynamics and solving for \dot{Q}_c , matched the measurements with a validation root-mean-square error of 2.3 kW .

4.3. Heating equipment model

4.3.1. Heat pump COP

In principle, the COP of the DC Nanogrid House’s variable-speed air-to-air heat pump depends on the indoor temperature and humidity, outdoor temperature, and compressor speed. In practice, modeling the COP as a function only of the outdoor temperature gave sufficient accuracy for control purposes. Fig. 7 shows how a quadratic COP model fits data from the manufacturer for varying indoor (vertical scatter) and outdoor temperatures. Although the model does not depend on

the indoor temperature or compressor speed, it fits the training data with an R^2 value of 0.98 .

4.3.2. Heating plant electric power

The heating plant consists of (a) the heat pump with electric power capacity (including the compressor and both fans) \bar{P} (kW) and COP η (a function of the outdoor temperature T_{out}), and (b) the heating elements with combined electric power capacity \bar{P}_r (kW) and unity COP. If the heat pump cannot keep up with heating demand, the device-level control system turns on the heating elements. This logic motivates modeling the electric power input to the heating plant as

$$P = \begin{cases} \dot{Q}_c/\eta & \text{if } \dot{Q}_c \leq \eta\bar{P} \\ \bar{P} + \dot{Q}_c - \eta\bar{P} & \text{otherwise.} \end{cases} \quad (10)$$

The above model involves a piecewise definition that does not easily integrate into optimization software without resorting to mixed-integer programming, which significantly increases computational complexity. By contrast, the equivalent electric power model

$$P = \frac{\dot{Q}_c}{\eta} + \left(1 - \frac{1}{\eta}\right) \max\{0, \dot{Q}_c - \eta\bar{P}\} \quad (11)$$

integrates directly with convex optimization software. This form shows that P is a convex function of \dot{Q}_c , since the right-hand side is the sum of a linear function and a convex function (the composition of the convex function $\max\{0, \cdot\}$ with the affine function $\cdot - \eta\bar{P}$). Integrating directly with convex optimization software guarantees that the MPC subproblems solve in polynomial time to global optimality. Fig. 8 illustrates the equivalence of the two electric power models.

4.4. Control algorithm

The supervisory control system uses MPC. At each time step, it reads the latest temperature and power measurements, downloads the latest weather forecast from an API service, predicts a trajectory of exogenous thermal powers, solves an open-loop optimal control problem to plan a trajectory of indoor air temperature set-points, and sends the first planned set-point to the device-level control system. The process repeats at the next time step.

This subsection describes the open-loop optimal control problem’s decision variables, objectives, constraints, and input data. Appendix 6.4

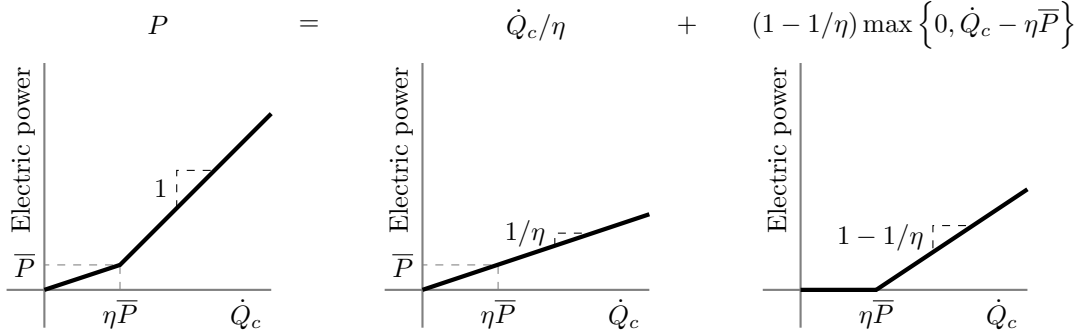


Figure 8: As a function of the HVAC thermal power \dot{Q}_c , the total electric power P (left) can be decomposed into the sum of a linear function with slope $1/\eta$ (center) and the positive part of a linear function with slope $1 - 1/\eta$ (right). This decomposition allows easy integration with convex optimization software, ensuring efficient global solution of the MPC subproblems.

describes several alternative problem formulations that the MPC system could use to pursue other objectives.

4.4.1. Decision variables

At each time k , the supervisory control system solves an open-loop optimal control problem that looks ahead over a receding horizon of L time steps. The $3L$ decision variables in this problem are trajectories over the prediction horizon of the indoor temperature set-points, HVAC thermal powers, and HVAC electric powers:

$$\begin{aligned} & T(k+1), \dots, T(k+L) \\ & \dot{Q}_c(k), \dots, \dot{Q}_c(k+L-1) \\ & P(k), \dots, P(k+L-1). \end{aligned} \quad (12)$$

After solving the optimization problem, the supervisory control system sends the optimal value of $T(k+1)$ to the thermostat as the indoor air temperature set-point for the next time step.

4.4.2. Objective function

The supervisory control system seeks to minimize the weighted sum of the peak electric power demand cost, the cumulative electrical energy cost, and the cumulative thermal discomfort penalty. The overall objective function is therefore

$$\begin{aligned} & \pi_d \max \{P(k), \dots, P(k+L-1)\} \\ & + \Delta t \sum_{\ell=0}^{L-1} \left[\pi_e(k+\ell) P(k+\ell) \right. \\ & \left. + \pi_t(k+\ell+1) |T(k+\ell+1) - \tilde{T}(k+\ell+1)| \right], \end{aligned} \quad (13)$$

where π_d (\$/kW) is the peak demand price, π_e (\$/kWh) is the electrical energy price, π_t (\$/°C/h) is the thermal discomfort price, and \tilde{T} (°C) is the occupant-specified reference indoor temperature. The control system balances trade-offs between the competing objectives by tuning the peak demand price π_d and the thermal discomfort price π_t . The last paragraph of this section discusses price tuning in more detail.

4.4.3. Constraints

The supervisory control system enforces the following constraints for each $\ell = 0, \dots, L-1$: the dynamics

$$\begin{aligned} T(k+\ell+1) = & aT(k+\ell) + (1-a)[\theta(k+\ell) \\ & + R(\dot{Q}_c(k+\ell) + \dot{Q}_e(k+\ell))], \end{aligned} \quad (14)$$

the heating plant model

$$\begin{aligned} P(k+\ell) = & \frac{\dot{Q}_c(k+\ell)}{\eta(k+\ell)} + \left(1 - \frac{1}{\eta(k+\ell)}\right) \max \left\{0, \right. \\ & \left. \dot{Q}_c(k+\ell) - \eta(k+\ell)\bar{P}\right\}, \end{aligned} \quad (15)$$

the heating capacity constraints

$$0 \leq \dot{Q}_c(k+\ell) \leq \eta(k+\ell)\bar{P} + \bar{P}_r, \quad (16)$$

and the thermal comfort constraints

$$|T(k+\ell+1) - \tilde{T}(k+\ell+1)| \leq \delta, \quad (17)$$

where δ (°C) is the maximum permitted absolute deviation of the indoor temperature from the occupant-specified reference temperature \tilde{T} . The time-variation of the COP $\eta(k+\ell)$ in Eq. 15 comes

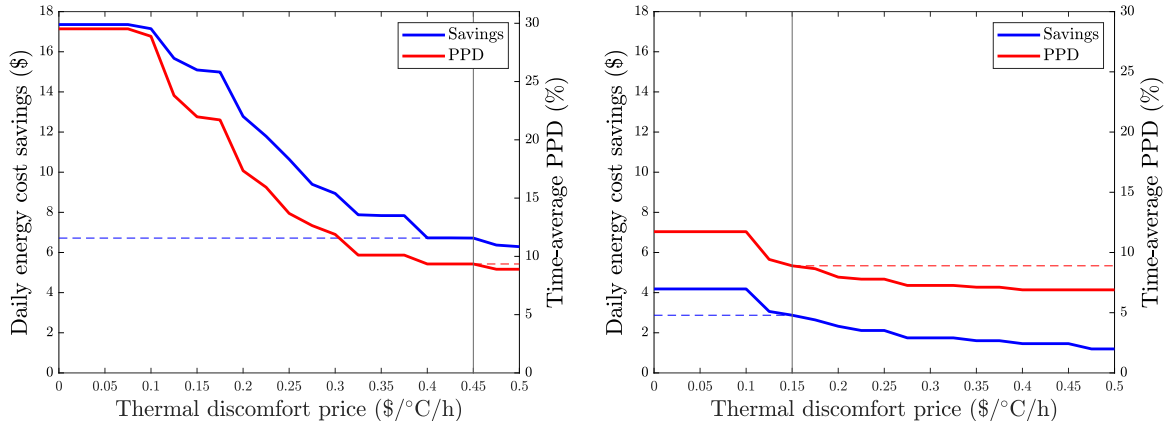


Figure 9: The tuned thermal discomfort price (vertical line) was higher on a cold day (left, average outdoor temperature -8.6 °C) than a mild day (right, 4.8 °C). On both days, the simulated time-average PPD was about 10%.

from the dependence of the COP on the outdoor temperature (shown in Fig. 7). This temperature dependence encourages the optimization to shift thermal load to times of higher COP.

4.4.4. Input data

In field demonstrations of the supervisory control system, we used time step duration $\Delta t = 1$ h, prediction horizon of $L = 24$ time steps, constant electrical energy price $\pi_e = 0.15$ \$/kWh, heat pump power limit $\bar{P} = 4.5$ kW, heating element capacity $\bar{P}_r = 19.2$ kW, temperature deviation $\delta = 3$ °C, and trained discrete-time dynamics parameter $a = 0.8$. The trained thermal resistances $R_{\text{out}} = 2.04$ °C/kW and $R_m = 1.06$ °C/kW generated the effective thermal resistance $R = 0.7$ °C/kW via Eq. 3, as well as the effective boundary temperature predictions θ via Eq. 2 with outdoor temperature predictions T_{out} from weather forecasts. We predicted the COPs η by propagating the outdoor temperature predictions through the quadratic COP model illustrated in Fig. 7. We predicted the exogenous thermal power Q_e by propagating weather predictions and time features through the support vector machine described in Section 4.2.4. The DC Nanogrid House occupants specified reference temperatures \tilde{T} of 20 °C during the day and 18 °C overnight.

4.4.5. Dynamic price tuning

We tuned the peak demand price π_d and thermal discomfort price π_t to balance the competing objectives of energy cost, peak demand, and thermal comfort. The utility serving the DC Nanogrid

House does not impose a peak demand charge. Rather, we included the demand price as an additional incentive to mitigate demand peaks by reducing use of backup resistance heat. We fixed π_d at $\$0.8$ per kW of daily peak demand, which corresponds to a monthly peak demand price of about 25 \$/kW, a typical value for monthly demand charges in commercial buildings in the USA. Fixing the peak demand price at a constant value reduced the price tuning problem to calibrating the discomfort price π_t . Tuning π_t interpolates between two extremes: for $\pi_t = 0$, there is no set-point tracking objective; while in the limit $\pi_t \rightarrow \infty$, the set-point is constrained to equal the user’s preference \tilde{T} . The supervisory control system automatically tuned the thermal discomfort price every 12 hours by sweeping an array of π_t values, solving an open-loop optimal control problem for each value, selecting the lowest π_t that maintained the time-average PPD below 10%,

$$\pi_t = \min\{\pi \mid \text{PPD at price } \pi \leq 10\%\}, \quad (18)$$

then multiplying the selected π_t by 1.1 during the day and 0.2 overnight. We found that this procedure promoted comfort during daytime hours while allowing greater energy savings overnight. The tuning procedure led to higher discomfort prices on colder days, as Fig. 9 illustrates.

5. Field demonstration results

This section summarizes field test results from January 30 through March 30, 2023. We tested MPC for 33 of these 59 days. We typically ran MPC

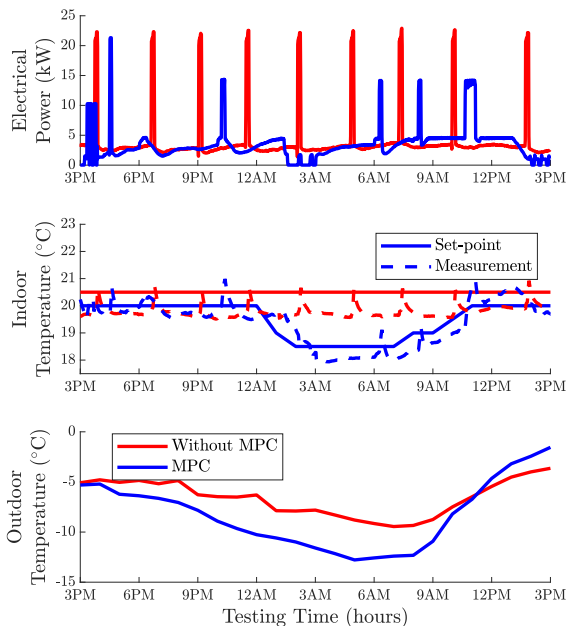


Figure 10: The heating equipment used 13% less energy on a cold day with MPC (blue) than a warmer day with constant temperature set-points (red).

for five-day stretches; the longest continuous MPC experiment lasted 12 days. We spent another 15 of the 59 days commissioning and debugging the control system. To benchmark performance, we compared MPC to the remaining 11 of the 59 test days and to historical data from past years.

We benchmarked MPC performance against the heat pump manufacturer’s default device-level control system with occupant-selected indoor temperature set-points and no supervisory control. The manufacturer does not publish the details of their device-level control system for competitive reasons. However, device-level control systems for variable-speed heat pumps typically use Proportional-Integral-Derivative loops that modulate compressor and fan speeds to track an indoor temperature set-point. Manufacturers typically wrap these loops within if-then logic that prevents the compressor from running below $\sim 30\%$ speed, limits ramping rates, implements minimum on- and off-times, and periodically defrosts the outdoor heat exchanger coils. This logic protects equipment and improves temperature regulation.

5.1. Control system behavior

Fig. 10 compares the control system behavior on a day with MPC (blue curves) to a day with a

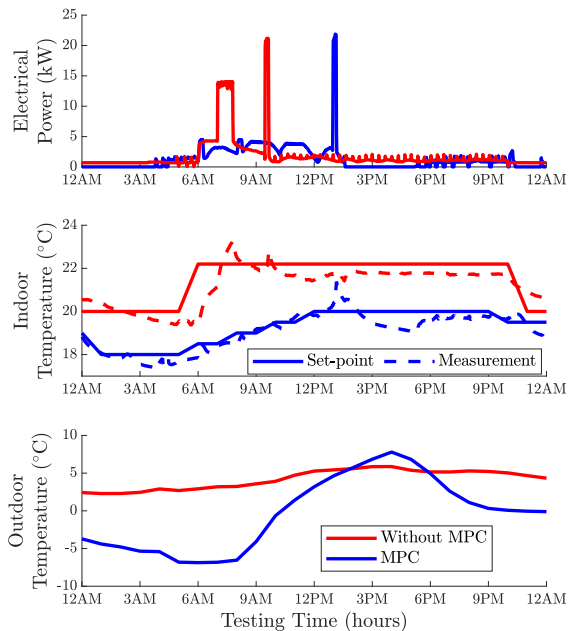


Figure 11: The heating equipment used 25% less energy with MPC on a cold day (blue) than on a warmer day with temperature set-point step changes at 6 AM and 11 PM (red).

constant indoor temperature set-point (red curves). On both days, the device-level control system kept the indoor temperature measurements (middle plot, dashed curves) reasonably close to the set-points (solid curves). On the MPC day, the outdoor temperature (bottom plot) reached a low of $-13\text{ }^\circ\text{C}$, compared to $-10\text{ }^\circ\text{C}$ on the constant-set-point day. Despite the colder weather on the MPC day, the heating plant used 84 kWh of electrical energy on the MPC day, 13% less than the 97 kWh used on the constant-set-point day. MPC reduced energy use mainly by reducing the indoor temperature overnight when the outdoor temperature dropped, causing the heat pump COP to drop (see Fig. 7) and heating demand to rise. MPC reduced the indoor temperature between midnight and three AM, then gradually warmed the house back up between six and nine AM. In addition to reducing energy, MPC reduced the frequency of using high-stage (19.2 kW) backup resistance heat from nine events on the constant-set-point day to one event on the MPC day. However, MPC increased the use of middle-stage (10 kW) resistance heat, primarily during the morning warm-up period.

It may appear from Fig. 10 that simply reducing the indoor temperature set-point overnight could yield similar energy savings to MPC. However, Fig.

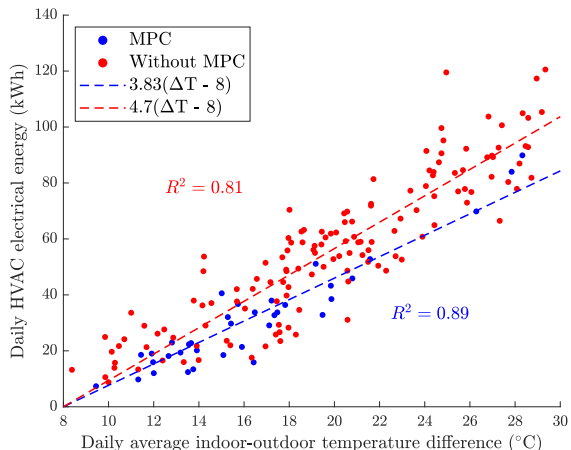


Figure 12: MPC typically saves about 20% of daily HVAC electrical energy. Absolute savings increase with the difference ΔT between the indoor and outdoor air temperatures.

11 shows that this is not the case in general. Fig. 11 compares an MPC experiment day (blue curves) to a non-MPC day (red curves) with two indoor temperature set-points: 22 °C during the day and 20 °C overnight. Although the MPC day was significantly colder (low of -7 °C) than the non-MPC day (2 °C), MPC used 27% less HVAC electrical energy than the 44 kWh used without MPC. The heating plant used more energy on the non-MPC day because the backup heating elements turned on in response to the sudden 2 °C set-point increase at six AM. While MPC also reduced set-points overnight, it warmed the house up more gradually in the morning, avoiding most use of the heating elements.

5.2. Energy

Fig. 12 shows the daily electrical HVAC energy with MPC (blue) and without MPC (red), as a function of the daily average temperature difference between the indoor and outdoor air. Energy use increases approximately linearly with the temperature difference both with and without MPC, but the slope is higher without MPC. Under the linear fits (dashed lines) in Fig. 12, the slopes are approximately Gaussian distributed. With MPC, the mean and standard deviation of the slope m are 3.83 and 0.117 kWh/°C. Without MPC, the mean and standard deviation of the slope \tilde{m} are 4.71 and 0.076 kWh/°C. The relative daily energy savings,

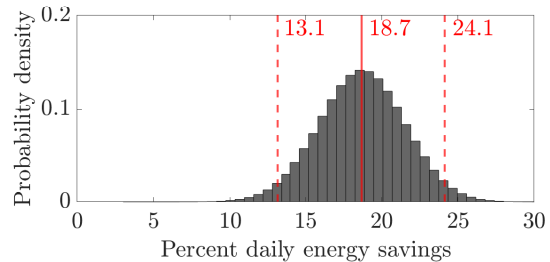


Figure 13: Histogram of the percent daily energy savings. Solid red line: mean. Dashed: 95% confidence interval.

as a fraction of the non-MPC daily energy use, are

$$1 - \frac{\text{MPC energy}}{\text{Non-MPC energy}} \approx 1 - \frac{m(\Delta T - 8)}{\tilde{m}(\Delta T - 8)} \quad (19)$$

$$= 1 - m/\tilde{m}.$$

Fig. 13 shows a histogram of the relative savings, based on 10^7 samples each from the distributions of m and \tilde{m} . The mean relative savings (solid red line) are 18.7%, with a 95% confidence interval (dashed red lines) of 13.1 to 24.1%.

Reduced use of backup heat contributed substantially to the total heating energy savings. On average over 19 days (4 with MPC, 15 without) with daily-average indoor-outdoor temperature differences near 22.5 °C, MPC reduced heating element energy use by 38%, from 15.7 kWh per day to 9.7. On average over these days, backup heat ran for 46 minutes per day with MPC and 55 without. During operation, backup heat tended to run at lower power with MPC than without.

Three effects can explain the observed reduction in backup heat use. First, MPC reduced heating loads by lowering indoor temperature set-points. This let the heat pump more often meet load without backup. Second, MPC sometimes preheated in anticipation of high loads that otherwise would have required backup heat. Third, lower loads during during defrost cycles (discussed further at the end of this section) caused the thermostat to request lower stages of backup heat.

The energy savings estimates presented here are somewhat conservative. A supervisory heating control system can save energy by reducing heating demand, operating equipment more efficiently, or both. Reducing heating demand entails reducing indoor temperatures. Our estimates do not account for energy saved by reducing heating demand as we normalized HVAC energy use by the indoor-outdoor temperature difference.

5.3. Thermal comfort

We assessed the supervisory control system’s thermal comfort performance in two ways. First, we periodically surveyed the two occupants using a digital survey. One or both occupants reported feeling uncomfortable on 4.8% of the 793 test hours with MPC, compared to 4.6% of the 724 test hours without MPC. Second, we used a thermal comfort software package [48] to calculate a PPD time series (see Section 4.4.5) based on indoor temperature measurements and the other inputs to the PPD model, such as occupants’ estimated clothing and activity levels. This process yielded a time-average PPD of 10.5% over the MPC test hours and 9.8% over the test hours without MPC. (Practitioners typically consider 10% PPD acceptable [33].) As the occupant comfort reports and PPD calculations produced similar metrics with and without MPC, we concluded that MPC maintained similar thermal comfort to the benchmark control system.

5.4. Demand peaks

In the DC Nanogrid House, peaks in electricity used for heating come primarily from the backup heating elements. Some use of heating elements is inevitable because the heat pump is not sized to meet heating demand without backup in the coldest weather. Furthermore, the device-level control system periodically reverses the flow of refrigerant (see Fig. 3), drawing heat from the indoor air to melt frost build-up on the outdoor heat exchanger [49]. The device-level control system turns on the heating elements during defrost cycles to avoid blowing cold air on occupants.

While no supervisory control system can completely avoid using the heating elements, MPC reduced the frequency and magnitude of heating element use. In the top plot in Fig. 10, for example, the frequency of heating element use dropped from nine events per day without MPC (red curve) to six with MPC (blue curve). In the six events where the heating elements turned on under MPC, they turned on at the 9.6 kW stage once, the 14.4 kW stage four times, and the 19.2 kW stage once. Without MPC, the heating elements turned on at 19.2 kW in all nine events. The electric power histograms in Fig. 14 show that these trends are not unique to the example days in Fig. 10. Across all field demonstrations, the heating elements ran at 19.2 kW during 48% of turn-on events without MPC and 8% with MPC (an 83% relative reduction). MPC reduced the mean heating element

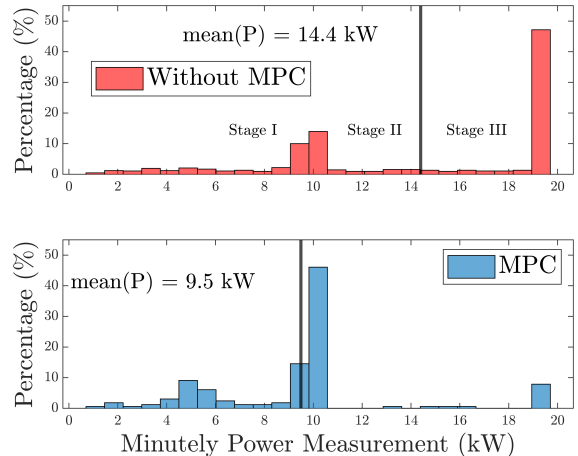


Figure 14: Histograms of the heating element power without MPC (top) and with MPC (bottom), conditioned on the event that the heating elements turned on.

power, conditioned on the event that the heating elements turned on, from 14.7 kW to 9.5 kW.

5.5. Defrost cycles

The device-level control system appeared to run fewer defrost cycles with MPC than without. On average over four MPC days with outdoor temperatures below zero °C and relative humidities between 70 and 80% (conditions under which frost often forms), the device-level control system ran 1.3 defrost cycles per day. By contrast, the device-level control system ran 2.7 defrost cycles per day over seven comparable days without MPC. We hypothesize that the reduction in defrost cycling came from MPC reducing heat demand in cold, humid weather. However, the sample sizes here are too small to draw robust conclusions. Future research could explore defrost effects in more detail.

5.6. Seasonal cost savings

The temperature-normalized energy models in Fig. 12 enable estimation of the total electricity cost savings from MPC in the DC Nanogrid House over the heating season of November 1, 2022, to March 31, 2023. The estimation procedure involves the daily average outdoor temperatures θ_i (°C) over each day $i = 1, \dots, 151$ of the heating season, the indoor temperature set-point \bar{T} (°C) without MPC, the difference γ (°C) between the daily average indoor temperature with and without MPC, the slopes m and \tilde{m} (kWh/°C) from Fig. 12, and

the electricity price $\pi_e = 0.15$ \$/kWh. Fig. 12 suggests modeling the total electrical energy used for heating on day i without MPC,

$$\tilde{E}_i = \tilde{m} \max\{0, \tilde{T} - \bar{\theta}_i - 8\}, \quad (20)$$

as zero if the daily average indoor-outdoor temperature difference $\tilde{T} - \theta_i$ is less than 8 °C and linear otherwise. The total electrical energy used for heating on day i with MPC follows a similar model:

$$E_i = m \max\{0, \tilde{T} - \gamma - \bar{\theta}_i - 8\}. \quad (21)$$

We estimated the seasonal electricity savings from MPC,

$$\tilde{E}_1 - E_1 + \dots + \tilde{E}_{151} - E_{151}, \quad (22)$$

using Monte Carlo simulation. The slopes m and \tilde{m} are normally distributed, as discussed in §5.2. Based on observations over the test days, we set $\tilde{T} = 20.7$ °C and modeled the daily average temperature reduction γ between default control and MPC as normally distributed with a 99% confidence interval of 0.7 to 1.7 °C. Over 10^6 Monte Carlo runs, these parameters gave a mean winter cost savings estimate of \$291 (95% confidence interval: \$239 to \$363). Relative to the non-MPC winter cost of \$1,029 (\$996 to \$1,061), MPC reduced winter heating costs by 28.3% (22.7 to 33.7%).

6. Practical challenges and scalability

Section 5 showed that MPC reduced energy costs and demand peaks while maintaining occupants’ thermal comfort. Optimistically extrapolated to all seasons and all residential buildings in the USA, which collectively spend about \$94 billion per year on space heating and cooling energy [50], the 23–34% heating energy cost savings shown here could translate to \$22–32 billion per year. While substantial, these potential cost savings alone do not justify deployment at scale. Implementation challenges, and the associated costs and labor requirements, also influence the economic case for deployment.

6.1. Practical problems and solutions

Deploying our supervisory control system in the field entailed solving a number of practical problems. Establishing communication between a diverse set of devices took time and persistence. In addition to the electric power sensing that our control system required, we installed various sensors

for validation purposes. These included flow and temperature sensors throughout the house’s thermal systems, which were nontrivial to install, network, and maintain. We recommend that future deployments rely only on electrical sensing and the thermostat’s temperature sensor. Pushing indoor temperature set-points to the variable-speed heat pump manufacturer’s proprietary thermostat also proved difficult. For future deployments, we recommend installing a thermostat with an open, user-friendly API. Several such thermostats work with most on/off equipment, although most variable-speed equipment today requires a manufacturer’s proprietary thermostat. Finally, we found that occupants occasionally felt uncomfortable regardless of the control scheme. We traced this issue to poorly balanced supply air distribution, which we fixed by adjusting damper positions on supply vents. Fixing this issue required trial and error over multiple in-person visits, a labor-intensive process that would not scale well.

6.2. Hardware and installation costs

Our field deployment required sensing the HVAC system’s electric power use and controlling its indoor temperature set-point. In principle, HVAC electric power can be estimated from whole-home measurements made by an electric utility’s smart power meter [51]. In practice, however, electric utilities rarely offer third-party access to real-time smart meter data. Therefore, replicating our field deployment at scale would likely require installing an electric power sensor for each new deployment. A suitable Internet-connected sensor that monitors power on eight circuits currently retails for \$125. Suitable Internet-connected thermostats currently retail for \$50–300. Most power sensor and thermostat installations involve skilled technicians. While costs of technician labor vary, installing a power sensor and thermostat would likely carry similar costs to the devices themselves. At a rough estimate, therefore, hardware and installation might cost several hundred dollars. A business’s actual hardware and installation costs would depend on their wholesale (rather than retail) device costs, and on their internal costs of technician labor. An HVAC manufacturer, for example, could outfit equipment with appropriate sensing and actuation hardware in the factory, rather than in the field. This could reduce hardware and installation costs.

6.3. Deployment labor

We spent approximately 190 engineer-days deploying, commissioning, and maintaining our supervisory control system. Of that labor, we spent about 150 engineer-days on tasks that we would not have to repeat if we deployed a similar system in another home. These one-time tasks included setting up a cloud database to store all data (low difficulty), writing code to pull data from sensors (medium), building a web form to survey occupant comfort (low), developing general methods to fit building and equipment models (medium), designing a software architecture for computing control actions (medium), developing a general control algorithm (medium), and establishing communication with the thermostat (high). We spent the remaining 40 engineer-days on tasks that we would have to repeat for each additional home. These repeated tasks included installing electric power sensors (0.5 engineer-days), fitting building and equipment models (20), tuning the control system (10), and maintaining the system despite network failures (10). As most team members had not deployed a supervisory HVAC control system before, we expect that labor requirements would decrease if we repeatedly deployed similar systems.

6.4. Limitations and future work

This work spanned one heating season. Running a supervisory HVAC control system for a full year would provide better estimates of annual cost savings. This work also showed an interesting interaction between the supervisory heat pump control system and device-level defrost controls. While it appeared that MPC reduced the rate of frost formation on the outdoor heat exchanger, our sample sizes were small. Longer and more controlled experiments, possibly building on recent frost modeling [52], could shed light on these effects. This work also involved little occupant participation in the supervisory HVAC control system. Future work could involve occupants more directly, for example through software that helps a user find their own balance between indoor temperatures, energy costs, greenhouse gas emissions, and impacts on the power grid.

CRediT authorship contribution statement

Elias Pergantis: Conceptualization, Methodology, Software, Investigation, Formal Analysis, Data

Curation, Visualization, Writing – Original Draft, Writing – Review & Editing. **Priyadarshan:** Methodology, Software, Data Curation, Writing – Original Draft. **Nadah Al Theeb:** Methodology, Data Curation, Formal Analysis, Writing – Original Draft. **Parveen Dhillon:** Methodology, Hardware, Data Analysis - Review, Writing – Review & Editing. **Jonathan Ore:** Software, Data Curation, Writing – Review & Editing. **Davide Zivani:** Conceptualization, Methodology, Formal Analysis, Review & Editing, Project Administration, Funding Acquisition. **Eckhard Groll:** Conceptualization, Methodology, Review & Editing, Project Administration, Funding Acquisition. **Kevin Kircher:** Conceptualization, Methodology, Formal Analysis, Writing - Review & Editing, Visualization, Project Administration, Funding Acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The publicly accessible repository [53] contains the weather and energy data used to train and validate the thermal models in this paper. The authors will make further data available upon request.

Acknowledgments

The Center for High Performance Buildings (CHPB) at Purdue University supported this work. The Onassis Foundation also supported E. Pergantis as one of its scholars. The authors would like to thank the occupants of the DC Nanogrid House for their patience and availability during the field demonstrations. Additionally, we extend many thanks to Trane Technologies for their technical support throughout the project.

Appendix A. Acronyms and notation

This paper used seven acronyms: HVAC (heating, ventilation, and air conditioning), MPC (model predictive control), RLC (reinforcement learning

Table 3: Mathematical notation

Symbol (Units)	Meaning
\bar{T} ($^{\circ}\text{C}$)	Indoor air temperature
T_m ($^{\circ}\text{C}$)	Thermal mass temperature
T_{out} ($^{\circ}\text{C}$)	Outdoor air temperature
\dot{Q}_c (kW)	HVAC thermal power
\dot{Q}_e (kW)	Exogenous thermal power
C (kWh/ $^{\circ}\text{C}$)	Thermal capacitance
R_m ($^{\circ}\text{C}/\text{kW}$)	Air-mass resistance
R_{out} ($^{\circ}\text{C}/\text{kW}$)	Indoor-outdoor resistance
θ ($^{\circ}\text{C}$)	Effective boundary temperature
R ($^{\circ}\text{C}/\text{kW}$)	Effective resistance
t (h)	Time
Δt (h)	Time step duration
a (-)	Discrete dynamics parameter
k, ℓ (-)	Discrete time indices
η (-)	HVAC system COP
P (kW)	Heat pump electric power
\bar{P} (kW)	Heat pump electrical capacity
\bar{P}_r (kW)	Heating element capacity
L (-)	Prediction horizon
π_e ($\$/\text{kWh}$)	Electricity price
π_d ($\$/\text{kW}$)	Peak demand price
π_t ($\$/^{\circ}\text{C}\text{h}$)	Thermal discomfort price
\tilde{T} ($^{\circ}\text{C}$)	Reference indoor temperature
δ ($^{\circ}\text{C}$)	Maximum temperature deviation
u ($^{\circ}\text{C}$)	Indoor air temperature set-point
τ (h)	Set-point tracking time constant
π_g ($\$/\text{kg}$)	Pollutant price
μ (kg/kWh)	Pollutant intensity
π_c ($\$/\text{kWh}$)	Curtailment price
\hat{P} (kW)	Demand response baseline
m (kWh/ $^{\circ}\text{C}$)	MPC energy savings slope
\tilde{m} (kWh/ $^{\circ}\text{C}$)	Non-MPC energy savings slope
$\bar{\theta}$ ($^{\circ}\text{C}$)	Mean outdoor temperature
γ ($^{\circ}\text{C}$)	Mean MPC temperature reduction

control), DC (direct current), API (application programming interface), COP (coefficient of performance), and PPD (predicted percentage dissatisfied). Table 3 summarizes the mathematical notation used in this paper.

Appendix B. Other problem formulations

The problem formulation described in §4.4 reflects the economic incentives and cyber-physical systems of the DC Nanogrid House. With minor modifications, this formulation could accommodate other settings. This section provides a few examples.

Other devices. The supervisory control system developed here could expand to include other

Internet-connected devices, such as electric vehicle chargers, home batteries, water heaters, or solar inverters. Adding one of these devices would require modeling the device and fitting the model parameters. By augmenting the open-loop optimal control problem with the new device’s model structure and parameters, the supervisory control system could operate the new device in concert with HVAC equipment to achieve house-level objectives, such as reducing peaks in the combined electricity demand of all devices.

Time-varying electricity prices. Although the DC Nanogrid House sees a constant electricity price, some residential customers face time-varying prices. Time-variation can manifest through Time-of-Use rates with two or more price tiers, hourly prices tied to wholesale market prices [54], or Critical Peak Pricing plans that raise prices when the grid is stressed. Given price predictions, the problem formulation described above can accommodate any time-varying pricing scheme.

Imperfect set-point tracking. The control algorithm in Section 4.4 implicitly assumes that the device-level control system accurately tracks indoor temperature set-points. We found this assumption accurate as long as set-points changed slower than ~ 2 $^{\circ}\text{C}/\text{h}$. The problem formulation could accommodate imperfect set-point tracking by specifying a tracking model, such as $dT/dt = (u - T)/\tau$, where u ($^{\circ}\text{C}$) is the indoor temperature set-point, T is the actual indoor temperature, and τ (h) is the tracking time constant; discretizing the tracking model in time; and appending it to the open-loop optimal control problem as a sequence of equality constraints. This formulation would add a trajectory of set-points as decision variables.

Pollutant emissions. The problem formulation could penalize emissions of greenhouse gases or other pollutants by adding the emission cost

$$\pi_g \Delta t \sum_{\ell=0}^{L-1} \mu(k + \ell) P(k + \ell) \quad (23)$$

to the objective function. This would require a pollutant price π_g ($\$/\text{kg}$) and forecasts of the pollutant intensity of electricity, μ (kg/kWh).

Grid-service revenues. Some retail demand response programs pay for electricity demand reductions below a baseline when the grid is stressed.

The problem formulation could include these incentives by subtracting the demand response revenue

$$\Delta t \sum_{\ell=0}^{L-1} \pi_c(k+\ell) \max \left\{ 0, \hat{P}(k+\ell) - P(k+\ell) \right\} \quad (24)$$

from the objective function. Here π_c (\$/kWh) is the curtailment price and \hat{P} (kW) is the demand response baseline. Similarly, some wholesale electricity markets allow demand-side participants to provide ancillary services such as regulation or reserve [55]. Reserve markets, for example, typically pay demand-side participants in day-ahead markets for committing to quickly ramp power down in response to generator failures or other contingencies. The problem formulation can accommodate wholesale ancillary service revenues through objective terms similar to those for demand response.

Changing formulations over time. The supervisory control system could readily adapt to new economic incentives or changes to cyber-physical systems. For example, if the occupants switched electricity plans, the control system could update the energy prices, peak demand penalties, grid-service incentives, *etc.* The control system could automate these adjustments based on input from occupants in a web dashboard or mobile application. The updated policy could go live immediately, eliminating any need to retrain a learning algorithm.

- [1] C. Delmastro, O. Chen, Buildings, Tech. rep., International Energy Agency (July 2023).
- [2] N. Sharma, A. Acharya, I. Jacob, S. Yamujala, V. Gupta, R. Bhakar, Major blackouts of the decade: Underlying causes, recommendations and arising challenges, in: 2021 9th IEEE International Conference on Power Systems (ICPS), IEEE, 2021, pp. 1–6.
- [3] K. Horowitz, Distribution system upgrade unit cost database current version, Tech. rep., National Renewable Energy Laboratory (2019).
- [4] U.S. Energy Information Administration, Electricity explained: Factors affecting electricity prices (August 2023).
- [5] E. Zhou, T. Mai, Electrification futures study: operational analysis of us power systems with increased electrification and demand-side flexibility, Tech. rep., National Renewable Energy Lab. (2021).
- [6] K. J. Kircher, K. M. Zhang, Heat purchase agreements could lower barriers to heat pump adoption, Applied Energy 286 (2021).
- [7] M. Killian, M. Kozek, Ten questions concerning model predictive control for energy efficient buildings, Building and Environment 105 (2016) 403–412.
- [8] J. Drgoña, J. Arroyo, I. C. Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, et al., All you need to know about model predictive control for buildings, Annual Reviews in Control 50 (2020) 190–232.
- [9] Z. Wang, T. Hong, Reinforcement learning for building controls: The opportunities and challenges, Applied Energy 269 (2020).
- [10] Z. Nagy, G. Henze, S. Dey, J. Arroyo, L. Helsen, X. Zhang, B. Chen, K. Amasyali, K. Kurte, A. Zamzam, et al., Ten questions concerning reinforcement learning for building energy management, Building and Environment (2023).
- [11] W. Stoecker, R. Crawford, S. Ikeda, W. H. Dolan, D. Leverenz, Reducing the peaks of internal air-conditioning loads by use of temperature swings, ASHRAE Transactions 87 (1981).
- [12] J. E. Braun, Methodologies for the design and control of central cooling plants, PhD thesis, The University of Wisconsin-Madison, 1988.
- [13] D. Spratt, Dynamic control – A case study, ASHRAE Transactions 95 (1989).
- [14] J. E. Braun, Reducing energy costs and peak electrical demand through optimal control of building thermal storage, ASHRAE Transactions 96 (2) (1990) 876–888.
- [15] G. P. Henze, R. H. Dodier, M. Krarti, Development of a predictive optimal controller for thermal energy storage systems, HVAC&R Research 3 (3) (1997) 233–264.
- [16] B. Kouvaritakis, M. Cannon, Model predictive control - classical, robust and stochastic, Switzerland: Springer International Publishing 38 (2016).
- [17] G. P. Henze, J. Schoenmann, Evaluation of reinforcement learning control for thermal energy storage systems, HVAC&R Research 9 (3) (2003) 259–275.
- [18] D. Blum, Z. Wang, C. Weyandt, D. Kim, M. Wetter, T. Hong, M. A. Piette, Field demonstration and implementation analysis of model predictive control in an office HVAC system, Applied Energy 318 (2022).
- [19] T. S. Pedersen, P. Andersen, K. M. Nielsen, Central control of heat pumps for smart grid purposes tested on single family houses, in: 2013 10th IEEE international conference on networking, sensing and control (ICNSC), IEEE, 2013, pp. 118–123.
- [20] B. Dong, K. P. Lam, A real-time model predictive control for building heating and cooling systems based on the occupancy behavior pattern detection and local weather forecasting, in: Building Simulation, Vol. 7, Springer, 2014, pp. 89–106.
- [21] A. Afram, F. Janabi-Sharifi, Supervisory model predictive controller (MPC) for residential HVAC systems: Implementation and experimentation on archetype sustainable house in Toronto, Energy and Buildings 154 (2017) 268–282.
- [22] F. Bünning, B. Huber, P. Heer, A. AbouDonia, J. Lygeros, Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings, Energy and Buildings 211 (2020).
- [23] C. Finck, R. Li, W. Zeiler, Optimal control of demand flexibility under real-time pricing for heating systems in buildings: A real-life demonstration, Applied Energy 263 (2020).
- [24] K. Kurte, J. Munk, O. Kotevska, K. Amasyali, R. Smith, E. McKee, Y. Du, B. Cui, T. Kuruganti, H. Zandi, Evaluating the adaptability of reinforcement learning based HVAC control for residential houses, Sustainability 12 (18) (2020).
- [25] M. D. Knudsen, L. Georges, K. S. Skeie, S. Petersen, Experimental test of a black-box economic model predictive control for residential space heating, Applied Energy 298 (2021).

- [26] F. Bünning, B. Huber, A. Schalbetter, A. Aboudonia, M. H. de Badyn, P. Heer, R. S. Smith, J. Lygeros, Physics-informed linear regression is competitive with two Machine Learning methods in residential building MPC, *Applied Energy* 310 (2022).
- [27] J. Vivian, L. Croci, A. Zarrella, Experimental tests on the performance of an economic model predictive control system in a lightweight building, *Applied Thermal Engineering* 213 (2022).
- [28] D. Wang, Y. Chen, W. Wang, C. Gao, Z. Wang, Field test of model predictive control in residential buildings for utility cost savings, *Energy and Buildings* 288 (2023).
- [29] S. Thorsteinsson, A. A. S. Kalae, P. Vogler-Finck, H. L. Stærmose, I. Katic, J. D. Bendtsen, Long-term experimental study of price responsive predictive control in a real occupied single-family house with heat pump, *Applied Energy* 347 (2023).
- [30] S. Brown, I. Beausoleil-Morrison, Long-term implementation of a model predictive controller for a hydronic floor heating and cooling system in a highly glazed house in Canada, *Applied Energy* 349 (2023).
- [31] D. Blum, K. Arendt, L. Rivalin, M. Piette, M. Wetter, C. Veje, Practical factors of envelope model setup and their effects on the performance of model predictive control for building heating, ventilating, and air conditioning systems, *Applied Energy* 236 (2019) 410–425.
- [32] J. Van Hoof, Forty years of Fanger’s model of thermal comfort: comfort for all?, *Indoor air* 18 (3) (2008) 182–201.
- [33] D. Enescu, A review of thermal comfort models and indicators for indoor environments, *Renewable and Sustainable Energy Reviews* 79 (2017) 1353–1379.
- [34] D. Sturzenegger, D. Gyalistras, M. Morari, R. S. Smith, Model predictive climate control of a Swiss office building: Implementation, results, and cost–benefit analysis, *IEEE Transactions on Control Systems Technology* 24 (1) (2015) 1–12.
- [35] M. Beccali, M. Bonomolo, F. Martorana, P. Catrini, A. Buscemi, Electrical hybrid heat pumps assisted by natural gas boilers: a review, *Applied Energy* 322 (2022).
- [36] A. Khakimova, A. Kusatayeva, A. Shamshimova, D. Sharipova, A. Bemporad, Y. Familiant, A. Shintemirov, V. Ten, M. Rubagotti, Optimal energy management of a small-size building via hybrid model predictive control, *Energy and Buildings* 140 (2017) 1–8.
- [37] L. Langer, T. Volling, An optimal home energy management system for modulating heat pumps and photovoltaic systems, *Applied Energy* 278 (2020).
- [38] J. Salpakari, P. Lund, Optimal and rule-based control strategies for energy flexibility in buildings with PV, *Applied Energy* 161 (2016) 425–436.
- [39] M. Yousefi, A. Hajizadeh, M. N. Soltani, B. Hredzak, Predictive home energy management system with photovoltaic array, heat pump, and plug-in electric vehicle, *IEEE Transactions on Industrial Informatics* 17 (1) (2020) 430–440.
- [40] J. Gasser, H. Cai, S. Karagiannopoulos, P. Heer, G. Hug, Predictive energy management of residential buildings while self-reporting flexibility envelope, *Applied Energy* 288 (2021).
- [41] I. Sharma, J. Dong, A. A. Malikopoulos, M. Street, J. Ostrowski, T. Kuruganti, R. Jackson, A modeling framework for optimal energy management of a residential building, *Energy and Buildings* 130 (2016) 55–63.
- [42] K. Zhang, A. Prakash, L. Paul, D. Blum, P. Alstone, J. Zoellick, R. Brown, M. Pritoni, Model predictive control for demand flexibility: Real-world operation of a commercial building with photovoltaic and battery systems, *Advances in Applied Energy* 7 (2022).
- [43] J. P. Ore, D. Ziviani, E. A. Groll, Evaluation of energy consumption and efficiency in a historical residential home before and after building renovations, *ASHRAE Transactions* 127 (1) (2021).
- [44] J. Ore, F. Meral, D. Ziviani, E. Groll, The case for DC: Motivation of modern topologies, DC-powered solutions, and applications within residential environments, in: *International High Performance Buildings Conference at Purdue*, 2021, pp. 1–10.
- [45] M. Grant, S. Boyd, CVX: Matlab software for disciplined convex programming, version 2.1, <http://cvxr.com/cvx> (Mar. 2014).
- [46] E. N. Pergantis, A. S. Sangamnerkar, Priyadarshan, J. Ore, P. Dhillon, D. Ziviani, E. A. Groll, K. J. Kircher, Sensors, storage, and algorithms for practical optimal controls in residential buildings, *ASHRAE Annual Conference*, Tampa, Florida (2023).
- [47] J. Ore, E. A. Groll, Design and development of a decentralized and distributed IoT home monitoring system within a DC nanogrid, in: *ASHRAE Topical Conference Proceedings*, American Society of Heating, Refrigeration and Air Conditioning Engineers, Inc., 2020, pp. 267–274.
- [48] F. Tartarini, S. Schiavon, pythermalcomfort: A python package for thermal comfort research, *SoftwareX* 12 (2020).
- [49] M. Song, S. Deng, C. Dang, N. Mao, Z. Wang, Review on improvement for air source heat pump units during frosting and defrosting, *Applied energy* 211 (2018) 1150–1170.
- [50] U.S. Energy Information Administration, Residential Energy Consumption Survey, Table CE3.6: Annual household site end-use expenditures in the United States – totals and averages (2020).
- [51] Z. E. Lee, K. M. Zhang, Scalable identification and control of residential heat pumps: A minimal hardware approach, *Applied Energy* 286 (2021).
- [52] J. Ma, D. Kim, J. E. Braun, W. T. Horton, Development and validation of a dynamic modeling framework for air-source heat pumps under cycling of frosting and reverse-cycle defrosting, *Energy* 272 (2023).
- [53] E. Pergantis, github.com/epergant/thermal_dataset (2023).
- [54] K. J. Kircher, K. M. Zhang, Model predictive control of thermal storage for demand response, in: *2015 American control conference (ACC)*, IEEE, 2015, pp. 956–961.
- [55] K. J. Kircher, Heat pump aggregation, optimization and control, PhD thesis, Cornell University, 2019.