Energy & Buildings 251 (2021) 111344

Contents lists available at ScienceDirect

Energy & Buildings

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

Development and evaluation of an automatic steam radiator control system for retrofitting legacy heating systems in existing buildings

Akram Syed Ali [a], Christopher Riley [a], Erica Acton [a], Amjad Ali [b], Mohammad Heidarinejad [a], Brent Stephens [a,*]

[a] Department of Civil, Architectural, and Environmental Engineering, Illinois Institute of Technology, Chicago, IL USA
[b] AKstudios, Schaumburg, IL USA

ARTICLE INFO

Article history:
Received 18 May 2021
Revised 30 July 2021
Accepted 10 August 2021
Available online 12 August 2021

Keywords:
Automatic control
Energy efficiency
Radiator heating
Sensing
Indoor environment

ABSTRACT

Inefficient building systems and a lack of building controls in existing buildings with legacy systems lead to increased energy consumption and reduced thermal comfort. Recent developments in low-cost sensors and microcontrollers have enabled the development and deployment of smart building sensing and controls solutions, but there remains a need for adapting these technologies to retrofit existing buildings with legacy systems. Here we introduce an automatic radiator control system, designed to address an application for which no commercial solutions currently exist, that can be used to retrofit low pressure steam heating systems in existing buildings. It allows for zone-level heat output control, which is adjusted continuously in real-time using data from custom wireless sensors. We describe the design and development of the physical, electrical, mechanical, and software components of the radiator control and present an evaluation of its performance while deployed in an historic building on the campus of Illinois Institute of Technology in Chicago, IL USA, over two winter seasons of 2019 and 2020. The automatic radiator control was deployed in a total of 17 rooms with varying uses and occupancy patterns. Five different control strategies were evaluated over a total of 15 weeks (9 weeks in 2019 and 6 weeks in 2020): (i) manual control (baseline), (ii) enforced schedule, (iii) occupancy-based control, (iv) proportional-integral-derivative (PID) control, and (v) PID and occupancy-based control. Each control strategy was applied one week at a time during a total of three different weeks, randomized throughout the experiment. Results show that using any of the automatic control strategies reduced radiator runtime compared to manual control, which reduces energy output. Using the PID + occupancy control strategy, which automatically adjusts the radiator heat output based on the setpoint alongside an occupancy sensor in each zone, yielded the highest potential energy savings during the 3-week-long measurement periods: approximately 74% compared to manual control. Using schedule- or occupancy-based control strategies led to some decreases in perceived comfort compared to manual control during the time of observation, but we also show that thermal comfort can be maintained and even improved by using a PID control strategy with the custom controller. In existing buildings where replacing legacy space heating systems with modern ones is financially and practically unfeasible, retrofitting them with custom automatic controls as demonstrated herein has the potential to considerably reduce energy consumption while maintaining or even improving thermal comfort, which can extend the life of the building.

1. Introduction

Studies have reported that energy consumption can be reduced in various building types by up to 30% by deploying sensors and automatic controls [1]. While modern buildings with building automation systems (BAS) use a combination of wired or wireless sensors, actuators, and building systems operating in several networks inside the building, an area of increasing interest is inexpensive wireless sensors and controls that can be used to create intelligent feedback loops in a wider variety of building types, including those without BAS. These systems have the potential to allow for better control of a wider variety of building systems such as mechanical systems (e.g., radiators, variable air volume (VAV) dampers, fans, pumps), lighting systems, enclosure systems (e.g.,

* Corresponding author at: Department of Civil, Architectural, and Environmental Engineering, Illinois Institute of Technology, Alumni Memorial Hall Room 228E, 3201 S Dearborn Street, Chicago, IL 60616, USA.

https://doi.org/10.1016/j.enbuild.2021.111344
0378-7788/

blinds), and other systems that impact energy consumption and thermal comfort. Wireless building controls based on open-source platforms such as Arduino and/or Raspberry Pi have been demonstrated in various buildings in the past several years [2–4]. These platforms have successfully been used to monitor indoor and outdoor environmental quality [5–8] and building energy use [9]. Increasingly, we now see highly customized control applications using such development platforms, such as controlling heat transfer for informing better heat exchanger designs [10] and demonstrating adaptive occupant-centered lighting control based on reinforcement learning [11].

Along with using open-source platforms, commercially available Wi-Fi-enabled thermostats also provide access to building operational data and allow for trend analysis and customized control. Large sample size studies using these newer connected thermostats have been conducted to measure residential HVAC system runtimes [12] and to geospatially track outages caused by major weather events such as hurricanes [13]. Scalable Wi-Fi-based home automation systems have also been implemented that allow for relatively low-cost, secure, and remotely controllable building control systems [14]. However, adoption of newer Wi-Fi enabled smart thermostats remains limited, both in terms of quantity and utility. For example, only 4% of homes have some kind of smart internet of things (IoT) device such as a smart thermostat [15–17]. While simple on–off control thermostats are widely used, programmable residential thermostats can sometimes use more energy in homes than those controlled manually, depending on how they are used [18]. One study showed that retrofitting packaged room-level heating, ventilation, and air-conditioning (HVAC) systems by implementing customized on–off controls augmented with environmental sensors around the room has the potential for substantial energy savings [19]. Thermostats used in commercial HVAC systems mostly operate on a fixed schedule [20], which can lead to an increase in energy usage [21]. Moreover, while thermostats are extensively used in new buildings and increasingly used in existing buildings, they have their limitations. For example, most commercially available thermostats also only measure air temperature around the device, whereas other environmental variables also influence thermal comfort, including radiant temperature, air speed, relative humidity, and solar heat gain [22,23]. Thermostats also most commonly operate only at the level of whole-building or multi-zone, often consisting of at least several rooms, which does not provide room-level control for individual occupants. This in turn prevents utilizing smart zoning strategies where the temperature in every room can be set at different levels according to the room geometry and orientation, as well as variable occupancy patterns. When smart zoning strategies are implemented, it becomes possible to save over 15% in energy consumption, while increasing thermal comfort levels by over 25% [24].

Over 80% of newly constructed buildings now have at least one of the following smart building technologies: artificial intelligence, mobile edge computing, 5G networks, renewable energy production and consumption, wireless sensor networks, and/or building automation systems [25]. However, implementation of these new technologies in existing buildings that operate on legacy building systems presents numerous challenges [26]. Since replacing these legacy systems on a large scale can be cost prohibitive, there is a need for innovation in developing inexpensive sensors and control technologies that can be easily retrofitted in existing buildings and that can allow legacy systems to become ‘smart’ systems, integrated in modern Building Management Systems (BMS). Implementation of wireless sensors and smart building controls, along with making energy usage data available to users in high granularity and allowing them immediate control of building systems, can produce a significant impact in reduction of energy usage in older buildings. One such type of legacy system that presents unique challenges for automatic control is hot water and steam radiator heating.

1.1. Challenges in controlling legacy hot water and steam radiator heating systems

The two main types of radiators used in the US are cast-iron radiators and baseboard style radiators that use aluminum-fin and copper pipes to distribute heat from hot water or steam [27]. While cast iron radiators are no longer used in new residential or commercial construction, a small number of specialty construction applications use ornate decorative cast iron radiators for aesthetics, luxury, and other historical reasons. A significant portion of the older existing building stock in the US uses hot water or steam radiators to meet their heating demands. Most buildings that use hot water or steam radiators in the US are concentrated around dense urban cities such as New York, Chicago, Boston, and Philadelphia, as well as in detached homes located in the Midwest and Northeast regions [28]. According to a report by Urban Green Council, around 80% of residential buildings are still heated by steam in New York City as of 2019 [29].

These legacy systems present opportunities and challenges for energy savings given that building space heating is the largest building energy end use in the US [15,30] and they are frequently oversized, leading to overheating and temperature imbalances among rooms (especially in multifamily housing) [31]. In some buildings with poor adaptation of control devices, the temperature can fluctuate by more than 25 °C [32]. Given these issues, there have been various efforts to reduce the heat output of these systems, which has mostly been accomplished by installing or redesigning the control valve at the zone level or by developing optimization and demand control heating models that adjust the heat output at the zone and/or building level.

The most common method of reducing energy use by radiator heating systems, especially in apartments and multifamily buildings, is installing thermostatic radiator valves (TRVs) [33]. These are self-regulating valves that typically contain a material such as a wax or a liquid that expands and contracts with changes in surrounding temperature, which then partially opens and closes the inlet valve of the radiator, consequently reducing or increasing the heat output. A study of 8 multifamily buildings in New York City showed that installing TRVs on radiators reduced average room temperatures and that total space heating energy use can be decreased by up to 16% [34]. One long-term field evaluation of using TRVs in residential buildings over several heating seasons in Lublin, Poland showed energy savings of up to 23%, with a payback time of less than 2.5 heating seasons [35]. TRVs on other types of heating systems apart from radiators such as floor heating systems also have the potential for energy savings of around 20–25% when used in conjunction with air and water temperature sensors for a feedback loop [36]. While some studies quantify thermal comfort by taking it into consideration, from the literature it is seen that building controls research tends to focus generally on achieving energy savings only, and do not always take into account the results of implementing the proposed control strategy on occupant thermal comfort and satisfaction [37].

There are also many studies showing the development of models for simulating the thermal and hydraulic behavior of space heating systems with radiators controlled by TRVs in multifamily buildings and their impacts on buildings [38,39]. Further, studies such as [40] show dynamic modelling of TRVs that aim to better understand the various dynamics behind their operation, such as the impact of hysteresis. While the goal of TRVs is to control the temperature of a room by reducing the heat output of the radiators, in practice they do not always provide optimal thermal

comfort. One reason for this is the material's slow expansion and contraction properties, which can result in users overcompensating by manually turning up the setpoint on the TRVs, and then adjusting down when it gets too hot. TRV users often complain that, having adjusted their valves to provide room temperatures to their liking in winter, further adjustments are needed to maintain the same comfort levels in milder weather [41]. It was also found that the TRV setpoints in individual rooms were higher in some cases than the whole house thermostat setpoint temperature, which implies that there may be a misunderstanding of the role of TRVs in residential heating systems [42].

Some additional efforts have been made to restrict the radiant heat being emitted by radiators by using a physical barrier such as an enclosure or a cover over the radiator [43]. One modern implementation of such studies include The Cozy™ by Radiator Labs, which is an internet-connected thermostatic cover for radiators that can sense the room temperature and control the heat output by using a fan to maintain optimal thermal comfort in the room [44]. Other commercially available products designed to control radiator heat output include automatic electronically actuated valve products (e.g., Honeywell HR90, Terrier 635011, Danfoss TWA, Netatmo NAV-EN). However, many of these existing products are available for purchase only in the European market, designed with thread sizes and materials (e.g., plastic) that are incompatible with many legacy high-temperature steam radiators in the US, as well as wireless radios that operate in a frequency spectrum that was allocated to the European market only. There are a few patents given to control devices for radiators in the European market [45,46], however they do not use wireless sensors, they do not have remote monitoring or control, and they are not designed for high-pressure steam radiators (which are especially prevalent in the U.S.). Most of these also operate on a set schedule, and do not have customizable control strategies based on zone type. This creates a need for an original radiator control device that can allow control of radiators and similar legacy heating systems, which not only take in input from wireless sensors, but also can have a fully customized control strategy specifically designed for a particular room to achieve both energy and thermal comfort goals.

1.2. Proposed legacy steam radiator control solution

To address the aforementioned limitations and drawbacks of existing radiator control devices such as TRVs and thermostats, as well as limitations in device communication in building automation systems [47], this study aims to develop and evaluate a custom control device for legacy steam radiator heating systems along with a fully featured backend system to allow real-time configuration of devices without the need to integrate into an existing BAS. Here we demonstrate the design and development of a customized steam radiator controller that communicates with multiple wireless sensors in individual rooms to better inform its controls and make smarter decisions for energy efficient operation. This custom solution operates on the *Elemental* platform, which our team recently developed as a combined wireless hardware and software platform that allows easy deployment and interoperability of wireless sensors and actuators, seamless integration of systems in existing buildings, availability of high-resolution real-time data for monitoring and control, and secure communication over the cloud where needed [48]. The custom radiator control hardware was developed from the ground up to fit legacy steam radiators in a historic building on the main campus of Illinois Institute of Technology but can be used in other similar systems. This customized radiator control provides the ability to control heat output at the individual zone level at each radiator. The goal of this project is to demonstrate the custom radiator control's performance deployed in the field and to demonstrate how it can be used to achieve potential energy savings while maximizing thermal comfort.

2. Methods

The following sections describe the design and development of the custom solution and the design of experiments and research methodology used to evaluate its performance.

2.1. Site description

The building chosen for this study was Alumni Memorial Hall, a historical Ludwig Mies van der Rohe academic building on the main campus of Illinois Institute of Technology in Chicago, IL USA. This building was constructed in 1945 and last renovated in 1972. Although BAS including Siemens direct digital control (DDC) or Delta Controls cover most of the buildings on campus, this building is not connected to any of these systems. It has a hydronic steam radiator heating system with two main radiator types: cast iron and aluminum fin tube radiators. Most of the radiators in the building are 2-column cast iron radiators of varying sizes, with some common areas having 6-column ones. Some rooms also have single or double stack aluminum fin radiators of varying sizes. The radiators operate on pressurized steam (between 34.5 and 68.9 kPa, or 5–10 psi), which is supplied throughout the building by a steam distribution network of pipes. All radiators have manual controls either in the form of TRVs or traditional gate valves. Additional details and schematics can be found in [28].

2.2. Existing radiator controls

Most of the radiators in Alumni Memorial Hall have Honeywell Braukmann thermostatic control actuators (model T104A1040, Fig. 1a). The radiators with gate valves were not considered in this project's scope due to complexity in adaptation and only those radiators with the thermostatic actuators were considered. The Honeywell actuators are specifically designed to control high-capacity radiators of pressures up to 100 kPa (15 psi) of steam and 1,000 kPa (150 psi) of hot water. They are operated manually by turning a dial and have guide marks between 0 and 6 to indicate the setpoint. The valves installed at each radiator are Honeywell Braukmann V110F1010 (shown in Fig. 1b), which have a ¾-inch thread designed specifically for the Honeywell thermostatic actuators. These valves are normally open without the control mounted. The valve pin is pushed by the actuator head located inside the actuator. This pin allows control over steam flow rate into the radiator. The actuator translates the rotational movement of the dial into an axial load by means of a thread, which in turn allows the valve to open and close. These products were widely used in the

Fig. 1. (a) Honeywell Braukmann thermostatic control actuators (model T104A1040) and (b) steam valve (Honeywell Braukmann V110F1010) controlled by the thermostatic actuator which can fit a typical steam or hot water radiator.

US at the time of the building's construction, so while our testing is limited to a single building, the controls can also be applied to other buildings with similar legacy system configurations.

2.3. Development of a custom radiator control system

The design and development of an automatic radiator control required specifically engineered mechanical and electrical components that work together to control the heat output in a room. The mechanical part of the design included modeling and fabricating custom 3D printed parts. The electrical part involved designing and assembling custom printed circuit boards (PCB) that work in tandem with the mechanical components and can be controlled wirelessly. Custom firmware and control logic for the radiator controls were also developed that allow real-time configuration and automation, and this was used to demonstrate potential energy savings. A fully featured backend software was also developed that allows remote configuration of each radiator control based on location, size of room, occupant preference and other variables. The radiator controls have custom Python scripts that allow independent operation, as well as allow local and remote access of controls for data collection and setting up customized control strategies. The custom radiator control is designed to integrate seamlessly with wireless sensors and the backhaul in the *Elemental* platform.

The physical design of the enclosure was approached with three considerations in mind: 1) mechanical conditions required to move the radiator valve (i.e., amount of torque required); 2) electrical components required for user interaction and control; and 3) physical specifications required for 3D printed components and to mount the assembled controller to the radiator valve. The following sections describe the various components of the radiator control with these considerations in more detail.

2.3.1. Mechanical design

Instead of recreating the functionality of the Honeywell manual actuator in a custom package, the decision was made to design and adapt around the existing actuators. These Honeywell actuators meet ASHRAE Standard 102–1989, which specifies methods of testing non-electric, non-pneumatic thermostatic radiator valves [49]. The valve and the actuator mounting thread have nickel-plated brass cast body and have components that are designed to be resilient in high-temperature and high-pressure applications. The advantage of developing a design based around the existing manual actuators allows for easy retrofitting and wide adaptation to various building system types when needed. Using the actuator manually (with no servo or controller attached) by twisting its head does not allow for automatic control of the room temperature. The rooms in this building frequently get too hot with steam constantly flowing, or they are not hot enough at a lower setpoint, requiring users to manually adjust the control to improve conditions. The current control is also slow to respond to immediate comfort needs, which leads to occupants adapting other methods to supplement the lack of speed, such as using an electric heater for instant heat or opening a window to quickly cool down the overheated room. This reduces productivity and increases energy use.

To allow quick and precise control over the actuator's motion, a carefully designed 'knob' is attached to it, which is controlled by means of a servo motor as shown in Fig. 2c. The actuator body has several grooves (Fig. 2a) which allow a carefully designed 'knob' to fit on. This 'knob' fills in the grooves with a tight friction fit and has strategically placed mounting holes for a high-power servo motor head to be mounted on (Fig. 2b). The servo was selected based on the minimum amount of torque required to completely rotate the dial on the actuator in both directions while it was mounted on the valve. This torque was measured using a torque tester on the actuator head and was found to be around 6.913 kgf·cm (6 lbsf·in). A 20 kgf·cm servo was selected to have adequate overhead and allow for smooth rotation on the entire range of motion. The servo also has a fully controllable rotational angle of 270°, operating speed of 0.16 sec/60° at 5 V power. The control angle of the servo which translates to the setpoint on the actuator is provided to the servo by digital commands over the servo's three-wire connector, which is connected to a control board as described in section 2.3.3.

Using a servo attached to the actuator allows controlling steam flow into the radiator by controlling the knob automatically, which in turn allows continuous control of the heat output of the radiator, thereby giving better regulation of the room temperature. Using wireless sensors at a location of interest in the room provides the controller with direct feedback on the room temperature, and the controller can be optimized for thermal comfort specifically in those regions of the room that are most used.

2.3.2. User interaction and experience

Since the actuator is now designed to be electrically controlled, it was important to implement proper user interaction. An initial survey of occupants in the building showed that most people never interacted with the manual actuator on the radiator because it had an industrial appearance, and they were concerned about damaging it. There is also no feedback mechanism on these actuators to indicate if an adjustment on the dial translated to a change in heat output. To facilitate better user interaction, the radiator control has an OLED screen for instant setpoint and room temperature feedback and buttons that enable an easier form of setpoint adjustment as needed, similar to a traditional thermostat that is widely used in most buildings today. All mechanical and electrical components are arranged to be assembled and connected together seamlessly. The entire contraption is enclosed within a carefully designed 3D printed enclosure as shown in Fig. 3.

2.3.3. Electrical circuit design

The servo, buttons, and the OLED screen all connect to a custom circuit board that facilitates all controls. The electrical circuit design of this 'control board' was done in Autodesk Eagle 9.1 (PCB design software). The board is designed to attach directly to a Raspberry Pi Zero W (v1.1). The control board along with the Raspberry Pi acts as the main control unit for all control logic to control the servo, implement control strategies, monitor temperatures, send and receive data and most importantly, control the valve on the radiator. The custom board has a servo control chip, along with a microprocessor and a radio to handle wireless data from wireless sensors around the room to inform the control of the status of the indoor environment.

Special care was taken to provide independent power supplies to the main control board and to the servo. The Raspberry Pi requires a stable 5 V power source to ensure that it runs smoothly. Since the servo can occasionally consume large amounts of current at the moment of actuation, it can drop the voltage down to 4 V, which can make the controller unstable and cause a reset. To solve this, a dual 5 V USB power supply was chosen as the main power source. This ensures that the 5 V power that goes to the Pi is independent of the 5 V power that goes to the servo.

The custom board has a control chip (NXP PCA9685) that can control the servo using PWM (Pulse-Width Modulation) to get precise adjustments of the valve. along with an OLED screen (128x32 pixels) and two buttons for the user to interact with. These along with other hardware components on the board can be seen in Fig. 4. The board also uses the same radio as the wireless sensors use (RFM69HCW) and the same semantics, making it directly compatible with the *Elemental* platform. This allows for easy real-time monitoring and analysis of data, and for customized control

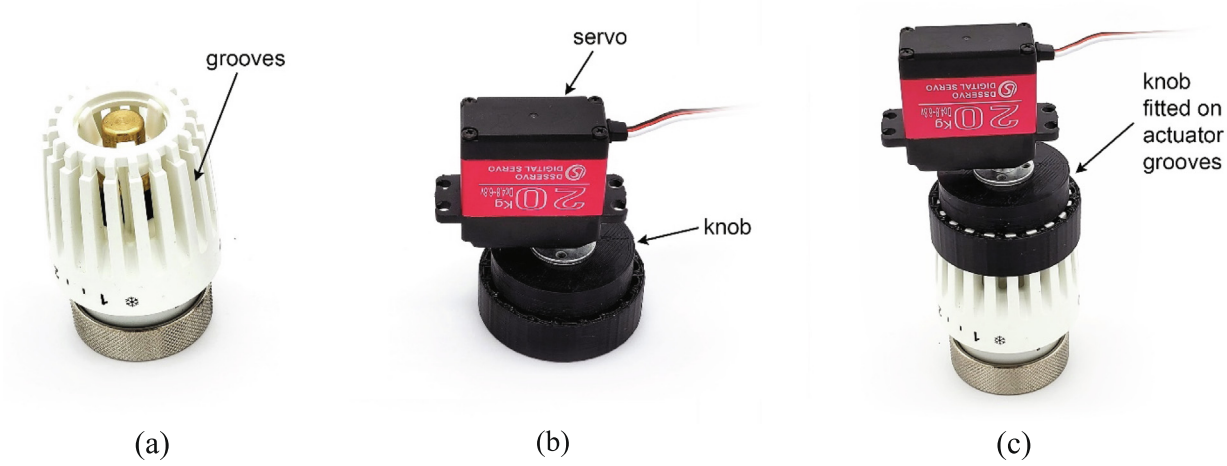

Fig. 2. (a) A thermostatic actuator with grooves, (b) a 3D printed 'knob' with a servo attached, and (c) knob + servo mounted on the actuator by filling in the grooves.

Fig. 3. A fully assembled automated radiator control node inside its 3D printed enclosure.

solutions for each room. An external temperature sensor on the board is also used, fixed directly on the radiator right after the valve, to detect the surface temperature of the radiator. This serves as an indicator of when the radiator valve is open and can also be used in fault detection in the control valves. Fig. 5a shows a control board attached to a Raspberry Pi Zero computer, and Fig. 5b shows a fully assembled control node in action while attached to a radiator valve and the dual USB power supply connected.

2.3.4. 3D printed components

Fig. 6 shows an exploded view of all components of the custom controller. The 3D printed components are all custom designed to fit around the mechanical and electrical components. The control cover acts as the main housing for all electronic components as well as the mount for the servo. It also hides all electronic components from the user, making the device more approachable, since the user only interacts with the buttons and the screen. The circuit board cover provides protection to the electrical components inside the controller. The buttons are designed to sit directly on top of the control board to allow easy actuation and tactile feedback. To ensure the rotational movement of the servo is properly translated to the actuator head and consequently the valve, the entire enclosure is bolted to a custom collar that was specifically designed to connect to the actuator mount as shown in Fig. 6. This mounts the assembled control node in place and allows full rotational motion from the servo to translate into the valve opening and closing. There are additional 3D printed spacers between the control board and the Raspberry Pi to prevent the boards from bending when the buttons are pressed.

All custom plastic components were 3D printed using a pair of Lulzbot Mini 3D printers and Lulzbot PolyLite PLA Polymaker filament. The collar was printed using colorFabb high-temperature PLA filament and additionally annealed in an oven at 110 °C for 30 min and then cooled down to increase its heat resistance capacity, since it would be in close proximity to the hot radiator valve.

2.3.5. Firmware on devices

The Raspberry Pi and the custom control board together form the "control node". The custom board's firmware is developed in Arduino (v1.8.5). Each control board is designed to receive data specifically from wireless nodes placed around in the room in which the control node is located. It ignores data from all other

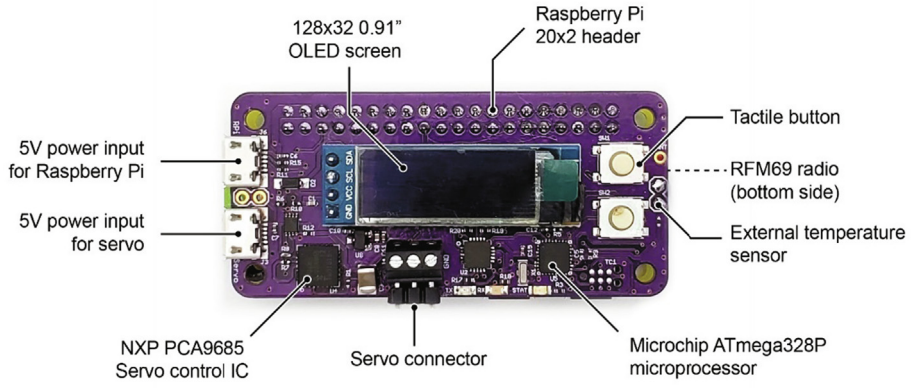

Fig. 4. Various components of a radiator valve control board.

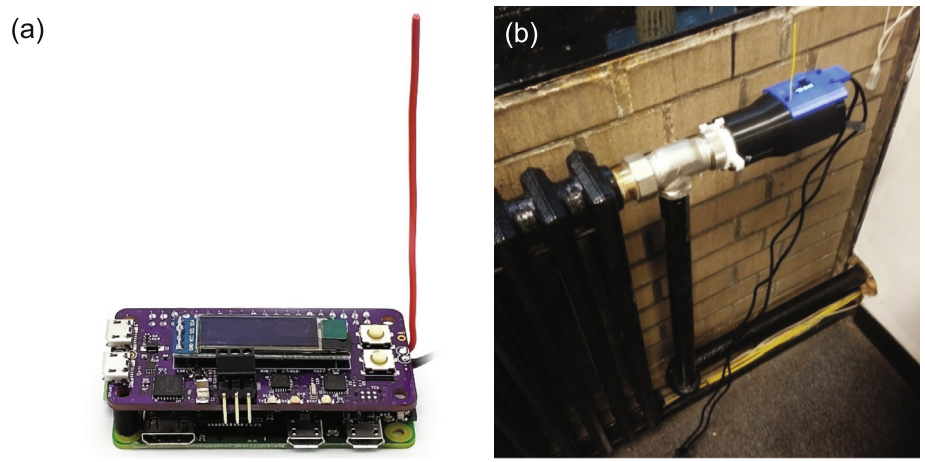

Fig. 5. (a) A control board with antenna attached to a Raspberry Pi and (b) a control node mounted on a radiator valve in the test building.

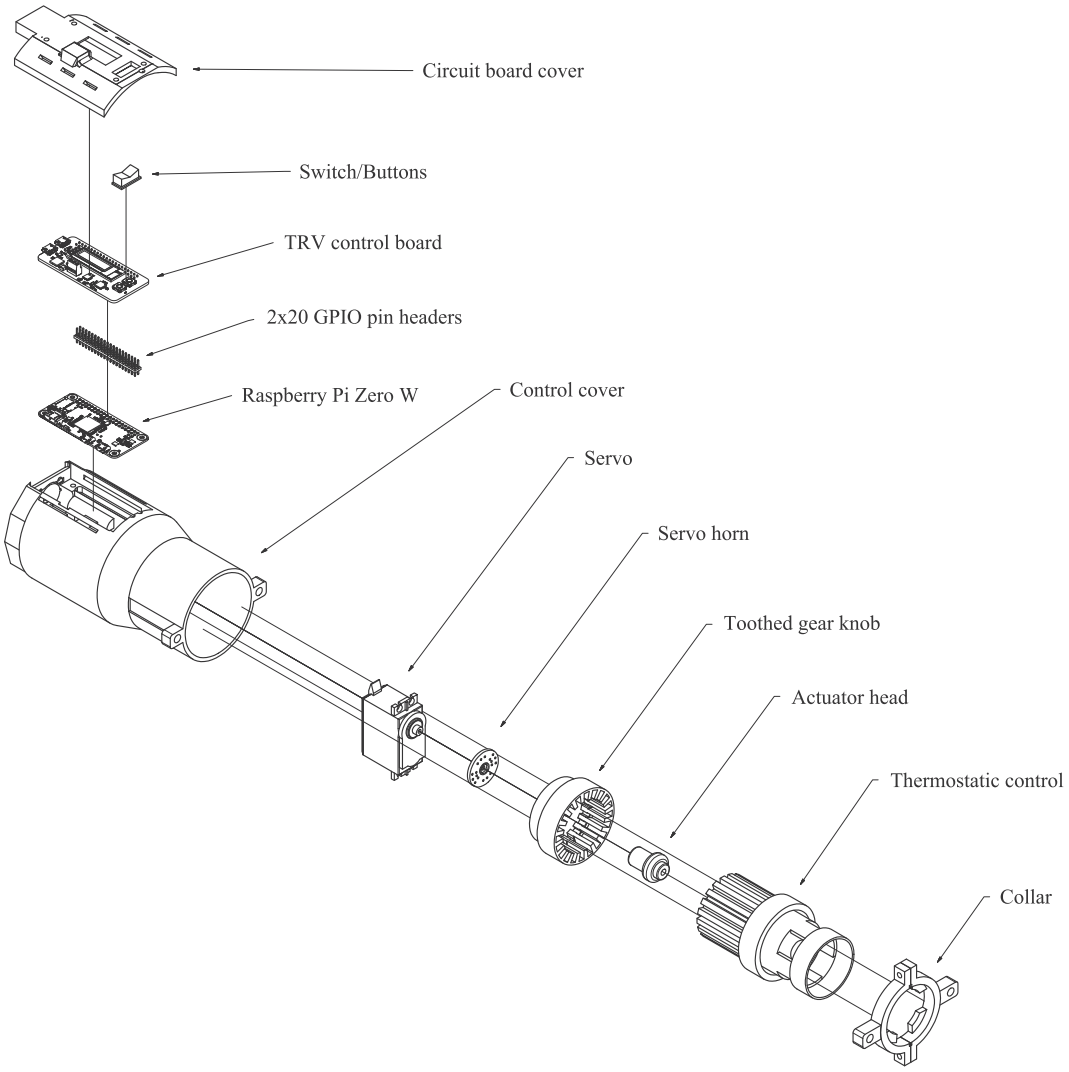

Fig. 6. An exploded view of all components of the control node.

wireless nodes in the building, to keep the control at the zone level and running independently of others. The board sends and receives data to and from the Raspberry Pi via serial communication. The wireless data are also sent back to the central backhaul that aggregates all the data and displays it in a web-based dashboard, which is described more in detail in Appendix A. The central backhaul can also be implemented directly on the Raspberry Pi, as explained in more detail in [48].

All control firmware on the Raspberry Pi is written in Python. A collection of interconnected Python scripts that work in tandem are responsible for operating the device. These scripts include:

1. *Configuration*: Periodically retrieve the latest configuration for a node via REST API and save it locally.
2. *Communication*: Handle all incoming and outgoing data over the serial port of the Raspberry Pi to and from the control board.
3. *User Interaction*: Handle button presses and controls what is displayed on the OLED screen for the user to see.
4. *Valve control*: Check the desired setpoint and adjusts the valve accordingly.
5. *Datalogger*: Log all data periodically to provide local backup in case the backhaul is inaccessible due to power loss or any other reason.
6. *Misc.*: There are other miscellaneous scripts running in the background that perform supplemental tasks, such as allowing for firmware updates, checking control strategies, overriding or adjusting any particular setting remotely, and internal diagnostics.

2.3.6. Radiator valve automation and control logic

The control nodes have custom firmware that implements various automatic and manual control strategies. This is in the form of four main Python scripts that are summarized as follows:

1. *Preheat*: This script implements a pre-defined preheat schedule for the radiator. Commercial buildings usually use a preheat schedule in early morning before occupants enter the building during cold months.
2. *Enforced Schedule*: This script implements an enforced schedule strategy for the radiators. This allows for the start and end time of the schedule along with a particular setpoint to be set so that the radiator operates within that given schedule and is forced to turn off outside those predefined hours. There is also a manual override implemented in this script which checks for manual user input after the enforced hours of operation and provides heating for a smaller additional time period.
3. *Check motion*: This script operates solely based on data from a motion sensor (if installed) in the room. The heat is available when the room is occupied, and the radiator is forced to turn off after a set period of inactivity in the room. This allows for aggressive energy savings, with a drawback of slightly reduced thermal comfort.
4. *Proportional-Integral-Derivative (PID) controller*: This script allows the control node on the radiator to operate like a traditional HVAC thermostat. It provides temperature-based setpoint to the user, and uses a custom PID control feedback loop, with real-time data coming from the temperature sensor in the room. The feedback-loop constantly adjusts the servo in precise increments to maintain the desired temperature setpoint. The PID control output is limited to a certain range and the number translates directly to the servo's position.

All automatic control strategies allow manual override by the user to take priority to ensure thermal comfort is maintained. Each control node is designed to get configuration settings from the backhaul so it can be remotely controlled and configured when needed. The configuration settings are constantly retrieved by a custom REST API that connects to a software running on the backhaul.

2.4. Deployment of devices

The automatic radiator controls were deployed in 17 rooms (19 radiators in total) around the test building. Fig. 7 shows the locations where the controls were deployed, which were spread across the first and second floor of the building. The gateway, which sends and receives all wireless sensor data, was deployed in a research lab located in the core of the building.

The controls had external surface temperature sensor probes that were attached to each radiator, which were used as indicators of radiator runtime. However, using the surface temperature as a surrogate for radiator usage does not take into account the time lag when steam starts to flow into the radiator and the room starting to warm up. It also cannot be used to predict energy usage accurately. To better understand the impact of this time lag and to calculate the actual energy output of the radiators, heat flux output from the surface of the radiators was measured at one-minute intervals. This was done using four pre-calibrated FluxTeq PHFS-01 heat flux sensors placed in equal spacing on the radiator, as shown in Fig. 8a. To make sure the entire surface area of the heat flux sensor is making contact with the radiator fin, thermal paste (Thermal Grizzly Kryonaut, with a thermal conductivity of 12.5 W/mK) was applied prior to installation (Fig. 8b). The sensor was attached to the radiator using 3M high temperature flue tape (Fig. 8c). All four heat flux sensors were connected to a FluxTeq COMPAQ DAQ and the serial output from this DAQ was sent to a Raspberry Pi, which logged all heat flux data locally on a microSD card.

Along with the radiator controls and heat flux sensors, over 95 custom wireless sensor nodes of various types (shown previously in [48]) were deployed around the building. Each room had a wireless temperature/RH sensor (Fig. 9a), motion sensor (Fig. 9b), and a mean radiant temperature (MRT) sensor (Fig. 9c). MRT is an important metric used in the calculation of thermal comfort using the Predicted Mean Vote (PMV) model [50], which is described in more detail in section 3.5. To estimate MRT, a custom hand-assembled black globe thermometer was placed in each room. This consisted of a temperature probe inserted into the center of a matte black ping-pong ball (Fig. 9c) connected to a wireless sensor, forming the MRT sensor.

According to ISO 7726, the black globe should be approximately 15 cm in diameter. If the globe is too small, the effects of air velocity are greater, reducing the accuracy of MRT measurement. However, the windows in the occupants' rooms do not open and the forced air HVAC system is largely unused during the cold winters in Chicago. Thus, the effects of natural and forced ventilation on varying ball diameter are negligible in this case since air velocity in the rooms is generally near zero. The mean radiant temperature is then calculated using Eq. (1) [51]:

$$MRT = \left[(GT + 273)^4 + \frac{1.110^8 v_a^{0.6}}{\varepsilon D^{0.4}} (GT - T_a) \right]^{1/4} - 273 \quad (1)$$

where MRT is the mean radiant temperature (°C), GT is the globe temperature (°C), v_a is the air velocity at the level of the globe (assumed to be 0.01 m/s), ε is the emissivity of the globe (assumed to be 0.95 for matte black plastic), D is the diameter of the globe (0.04 m for the ping-pong ball) and T_a is air temperature (°C).

2.5. Design of experiments

A design of experiments was developed to establish the energy saving potential of the controls. Five different control strategies were established to test practical operation in buildings with a defined occupancy schedule, such as commercial and academic buildings. These control strategies were:

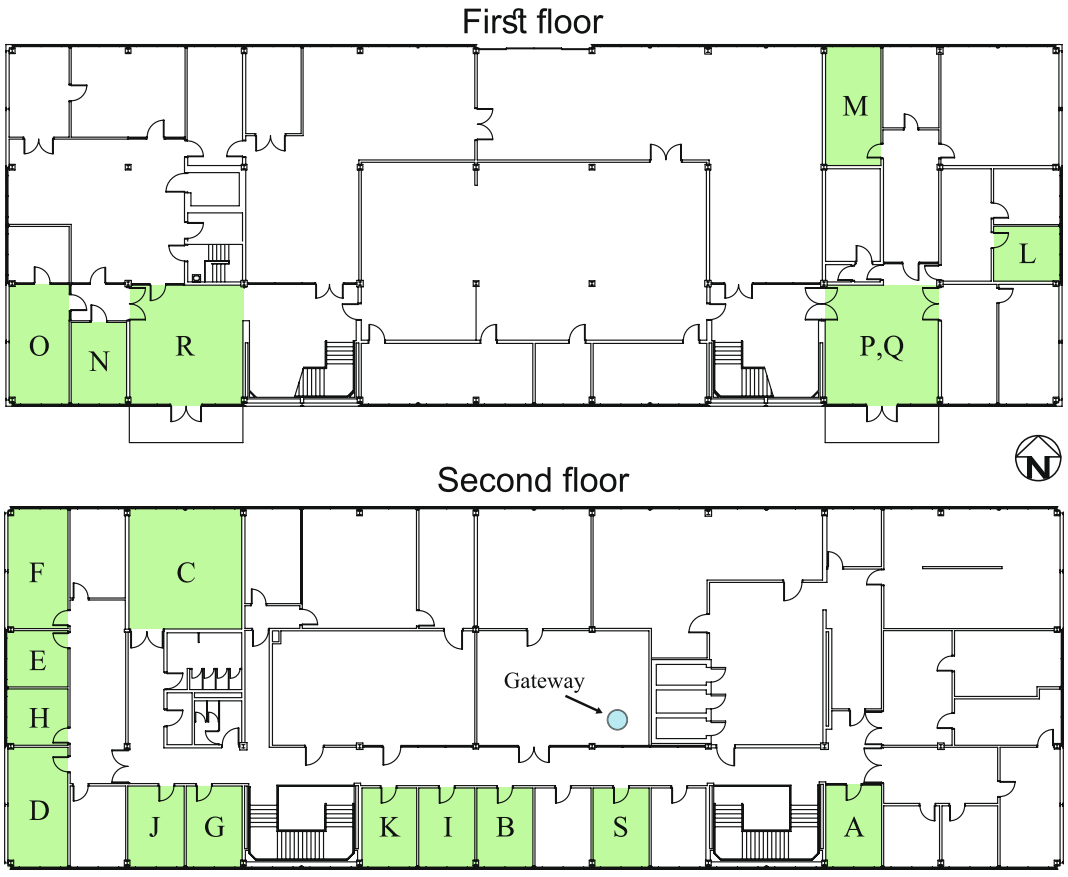

Fig. 7. Location of controls and gateway deployed in first and second floor of Alumni Hall.

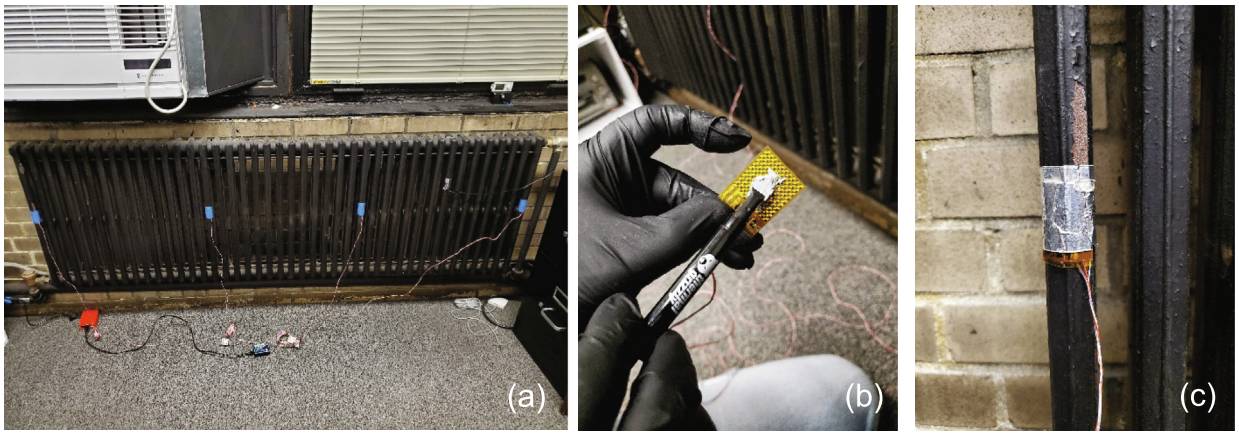

Fig. 8. (a) Heat flux sensors installed on radiator, (b) thermal paste on sensor to help in conductivity and (c) high temperature tape to attach sensor to radiator.

A. Manual control – Allow full manual control so the user can adjust the desired setpoint at will. This is similar to using the manual actuator head, but with a screen and buttons instead. The setpoints shown on the screen correlate to the setpoints on the actuator head.
B. Enforced Schedule – A start time and end time is set that automatically turns the radiators on and off on a fixed schedule, e.g., 7AM – 5PM. The user can override the setting past this time period for a short duration when extra heating is needed.
C. Occupancy-based control – This control strategy makes the radiator operate only when the room is occupied, and the radiator is forced to turn off after a set period of inactivity in the room (30 min). This allows for aggressive energy savings, with a drawback of slightly reduced thermal comfort.
D. PID control – This control strategy is similar to the manual control, except it allows the user to set a desired temperature instead of a valve setpoint. The radiator is set to operate like a traditional HVAC thermostat that uses PID control [52]. It provides temperature-based setpoint to the user,

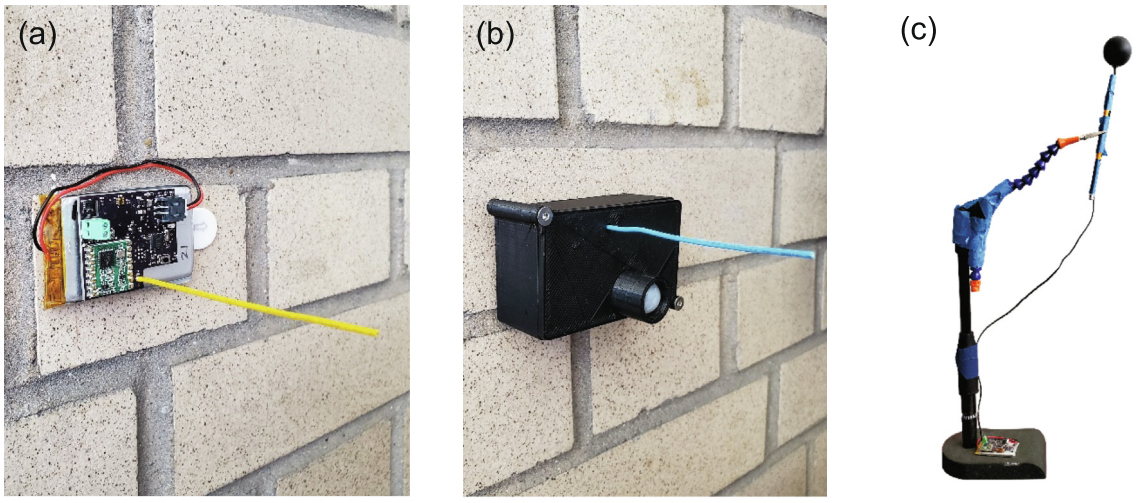

Fig. 9. (a) wireless temperature/RH sensor node, (b) wireless motion sensor node, and (c) a custom-built black-globe thermometer.

and uses a custom PID control, with data coming from the wireless temperature sensor in the room. The feedback-loop constantly adjusts the servo in precise increments to maintain the desired temperature setpoint.

E. PID + Occupancy-based control – This control strategy combines PID and occupancy-based control to allow room temperature control when needed and energy savings when the room is unoccupied.

Control strategies A, B, and C were implemented in the winter of 2019, while control strategies D and E were implemented in the winter of 2020. The control strategies for each control were randomly changed every week. The randomization was done to observe occupant interaction with the control, to assess thermal comfort for each strategy and to get some variation in outdoor weather during all control strategies. Twelve rooms were single-occupant faculty and staff offices, while 5 rooms were multi-occupant rooms used by students, research assistants, and teaching assistants. The controls were also deployed in a conference room and a classroom that were occasionally used by a larger group of people. The occupants were given a demonstration of what the control nodes are designed to do, and how to interact with them. They were also informed about the various control strategies but were not made aware of which strategy was used each week so as not to influence the results. A survey of room occupants (based on [53]) was also taken each week to assess thermal comfort. The survey included inquiries about occupant thermal sensation and satisfaction at the time of survey (mid-day) and also how occupant's felt in the morning that day when they first arrived in the room.

3. Results

The custom radiator control nodes were deployed in Alumni Memorial Hall (Chicago, IL) for the winter seasons in 2019 and 2020. The radiator control experiment ran for a total of 15 weeks (9 weeks in 2019 and 6 weeks in 2020). Each control strategy was operated for one week at a time during three different weeks randomized throughout the experiment, resulting in a total of three weeks of deployment for each control strategy across the two years. Each user had the option to manually override their radiator control outside of these control strategies, if they wanted to adjust the heat as needed. Fig. 10 shows a timeline of all control strategies for both winter seasons.

3.1. Radiator runtime

Fig. 11 shows the average radiator usage for each control strategy for the entire 2-year duration from a sample of the rooms that had controllers and sensors installed. Radiator usage was calculated based on the radiator surface temperature. It was observed that when the radiator valve was fully open, the maximum temperature on the surface of the radiator reached between 85 and 95 °C and dropped down to ambient air temperatures (20–22 °C) when the radiator was fully turned off and no steam was flowing through the system. Therefore, the radiator was assumed to be in-use when the radiator surface temperature increased above 40 °C, starting from room temperature, and was flagged in the data accordingly to calculate system runtimes.

From Fig. 11, there is a noticeable decrease in the amount of time the radiator was operating when automatic control strategies were used compared to manual control, with enforced schedule at an average of about 44% and occupancy-based control at an average of around 23% radiator usage, compared to manual control at an average of 61%. The data also showed that most occupants generally did not turn the radiators off when leaving the room. Setting the control strategy to PID control allowed for a slight reduction in radiator usage from manual control to an average of about 49%, as the PID control opens and closes the valve on the radiator to maintain a given room temperature, while the manual control keeps the valve open and a particular set point. Adding occupancy detection to PID control allowed for greater reductions in radiator usage, with the radiator running only about 24% of the time, on average.

It is also worth noting that a small number of rooms had high radiator usage even when automated control strategies were used. These anomalies were identified as being caused by faulty steam valves on those individual radiators that did not fully close and consequently leaked steam continuously, providing heat during all periods of the day and night even if occupants set their manual controls to the 'off' position. Identifying this fault allowed the maintenance team to immediately rectify the issue and save energy from being wasted further. The display dashboard has alerting features that can be customized to send a text or an email if certain criteria is met, such as a high radiator or room temperature observed when the radiator is supposed to be turned off.

To better understand the impact of outdoor weather on radiator output with the different control strategies, runtime data were normalized against outdoor temperature for the time period of the experiment. The weather data were extracted from

	Tes conditions	1	2	3	4	5	6	7	8	9											
A	Manual	A				A			A												
B	Enforced Schedule		B	B				B													
C	Occupancy-based				C		C			C											
D	PID																	D		D	D
E	PID + Occupancy																		E		
	Date	1/14/2019	1/21/2019	1/28/2019	2/4/2019	2/11/2019	2/18/2019	2/25/2019	3/4/2019	3/11/2019	3/18/2019	..	..	1/27/2020	2/3/2020	2/10/2020	2/17/2020	2/24/2020	3/2/2020	3/9/2020
		9-week experiment sart								9-week experiment end					6-week experiment sart				6-week experiment end		
		2019													2020						

Fig. 10. Timeline of control strategies. Each numeric value in the header refers to a weeklong test period.

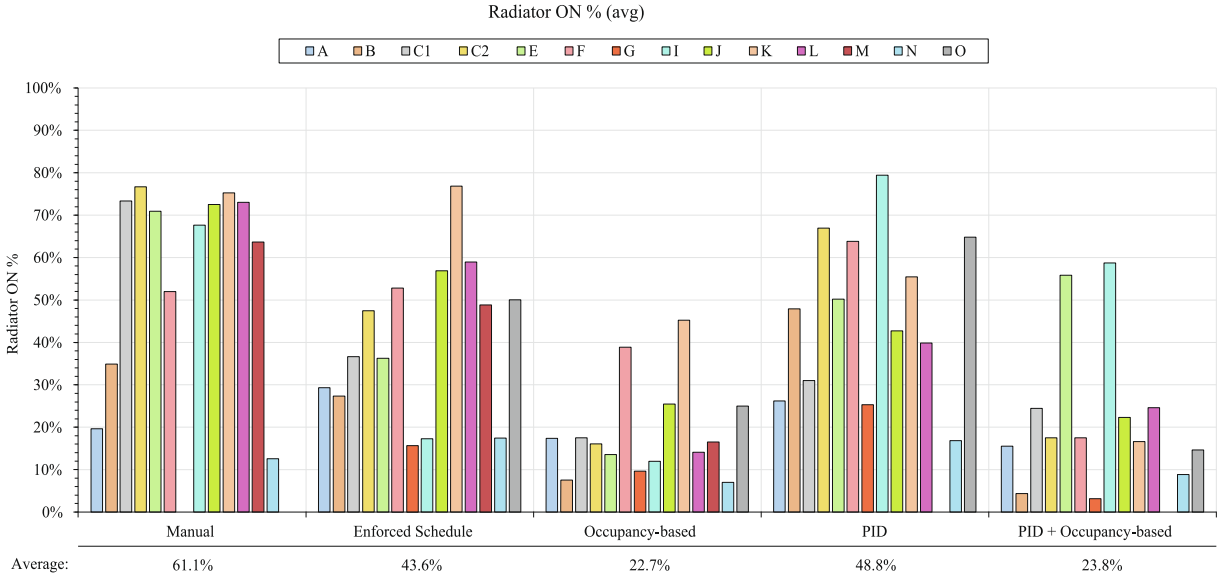

Fig. 11. Average radiator usage in Rooms A through O, expressed as percentage, for each control strategy for the duration of the experiment (both 2019 and 2020 heating seasons).

Wunderground [54], with the source of the data selected to be the weather station at IIT tower located at 10 W 35th St, Chicago, IL 60616. Fig. 12a and 12b show the daily average radiator operating usage for each room expressed as a percentage against the daily average outdoor temperature. The grey bands indicate the 95th confidence interval of the regression fit.

From Fig. 12, there is a generally downward trend in radiator usage for all automatic control strategies when outdoor temperature increases, while we observe almost no correlation when manual control was used. This indicates that occupants almost never manually turned the radiator off when leaving the room for the day. During warmer weather, we can see that there is a large potential for energy savings by implementing any kind of automated control that turns the radiators off when not needed. The two outliers seen on the leftmost side in Fig. 12a for the Enforced Schedule control strategy occurred during some of the

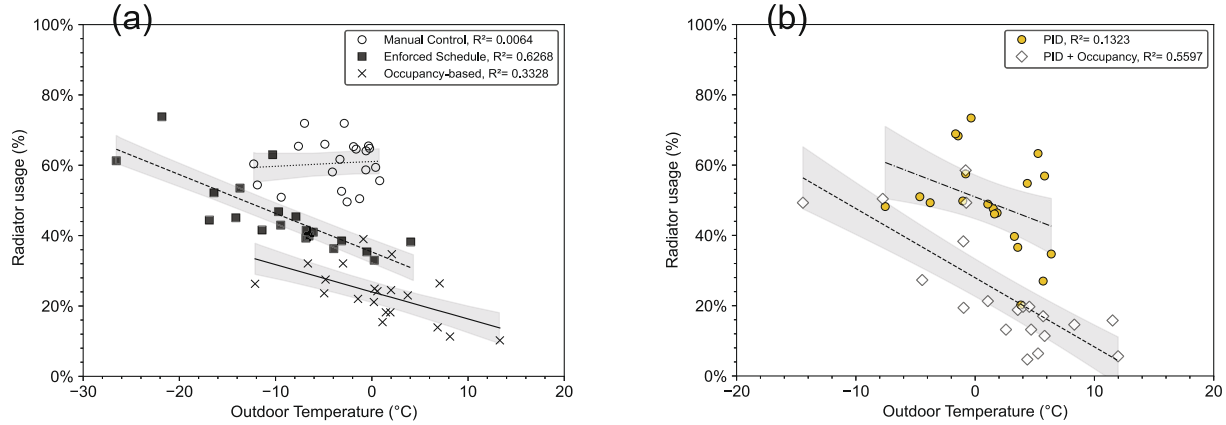

Fig. 12. Daily average radiator usage for each control strategy versus outdoor temperature for (a) manual, enforced schedule and occupancy-based; and (b) for PID and PID + occupancy-based.

coldest days of Chicago in the winter of 2019 (during a 'Polar Vortex' that season). While this skewed the data, rest of the data show similar trends in decreasing radiator usage with increasing outdoor temperature. Similar downward trends were also observed for both PID control strategies. However, the winter of 2020 was warmer than usual in Chicago, leading to lower radiator usage than normal.

3.2. Room occupancy patterns

To better understand occupancy patterns in each room, Fig. 13 shows an occupancy heat map for all rooms, generated using data from motion sensors, aggregated by the average frequency of motion events during every hour throughout the experiment. The frequency varies from 0 to 100%, which indicates how frequently the motion sensor detected occupancy in the room in that hour in the duration of the experiment.

From the occupancy heat map, we see that while some rooms show an occupancy pattern of a typical workday (9 AM to 5 PM), many others are occupied for a longer duration throughout the day. For example, Room C is a classroom that showed increased periods of occupancy at fixed times when classes were scheduled but was also occupied throughout the day on certain days by students who use it to study after hours. Rooms P, Q, and R are open lobby areas with seating located at the entrance of the building. These areas were also occupied throughout the day by students using the seating for leisure or study. The rest of the rooms were primarily single-occupancy offices, which displayed occupancy patterns that varied based on the occupant and their schedules. These data can be useful in future prediction of occupancy patterns based on day of the week and can allow better preheating to satisfy thermal comfort as well as custom heat scheduling at the zone level instead of the building or system level, although this is not explored here.

In order to understand the variation between occupancy patterns and control strategies, Fig. 14 shows distributions of the average daily room occupancy for all rooms during the entire duration of the experiment. We can see that the rooms had a median daily average occupancy of about 18% during all control strategies, with averages ranging from ~ 16% to ~ 22%. The enforced schedule control strategy showed the highest variation in daily average occupancy. The outliers in the graph are the lobby areas where the average occupancy tended to be much higher during the day as compared to all other rooms. These data show that although there was some natural variability in occupancy while different control strategies were implemented, the differences were relatively small and therefore we can make reasonable comparisons among the different control strategies without needing to control for occupancy.

3.3. Occupant engagement with radiator controls

To better understand how occupants interacted with the radiator controls, all button presses on the control to adjust the setpoints were logged, along with when a button was pressed to change a setpoint set during automated control strategies. This provides an additional surrogate for thermal comfort of the occupants. Fig. 15 shows plots of manual overrides of setpoints for all radiator controls during each week of the experiment. The plots on the left indicate when the occupants are lowering the setpoints and the plots on the right are manual overrides when the occupants were increasing the setpoint. It was seen that the occupants usually turned the heat output down more frequently than they turned it up, especially during warmer weather. This can be explained by the close proximity of most occupants to the radiators themselves, the large size of each radiator, and the presence

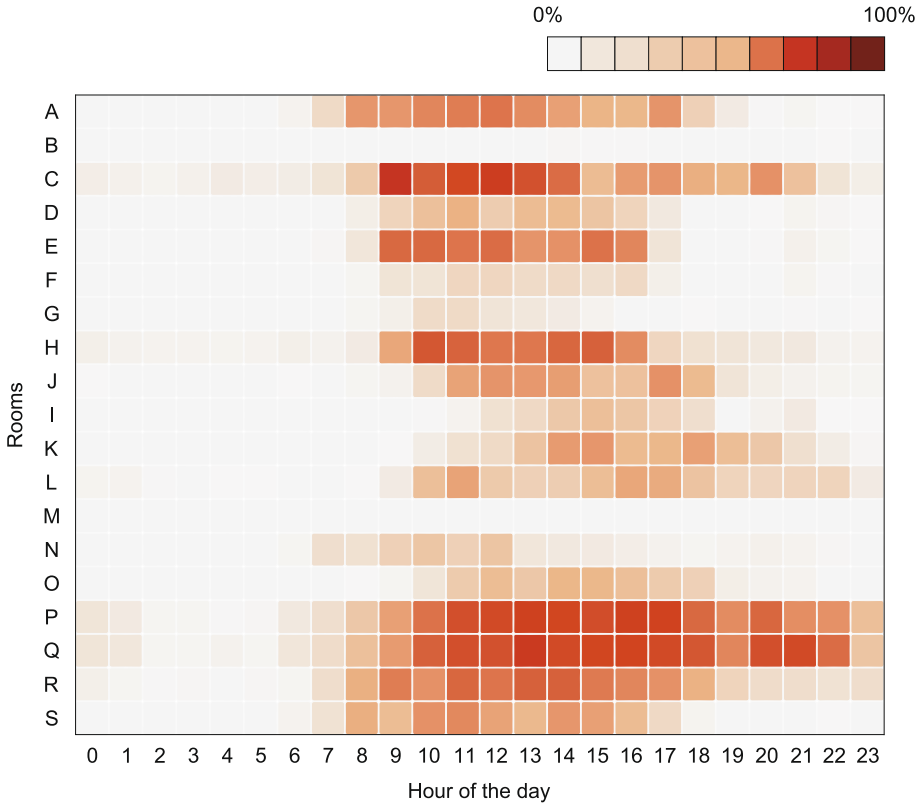

Fig. 13. Hourly occupancy heat map for all rooms.

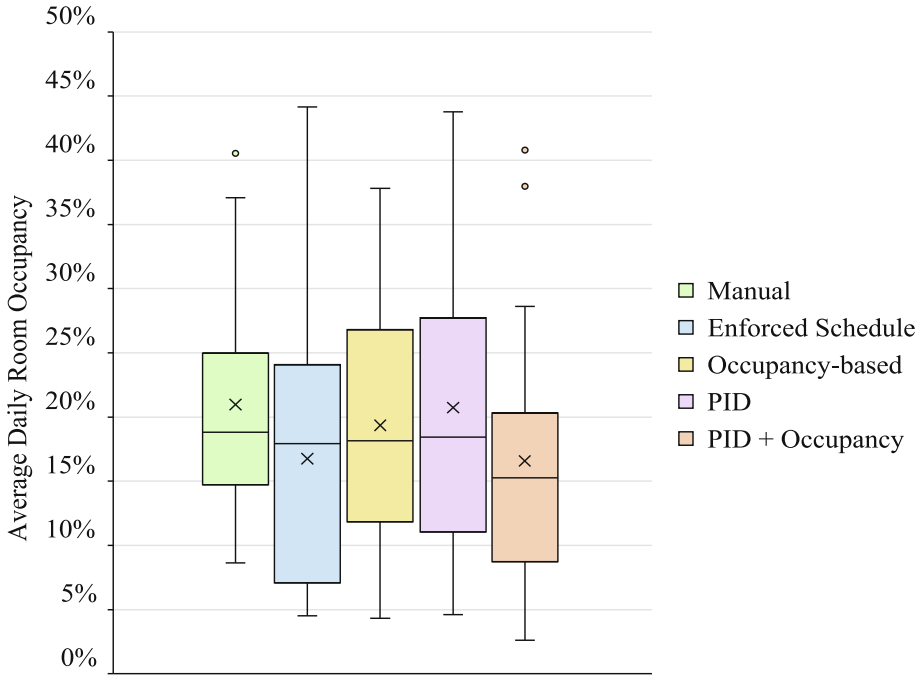

Fig. 14. Average daily room occupancy for all control strategies during the entire duration of the project.

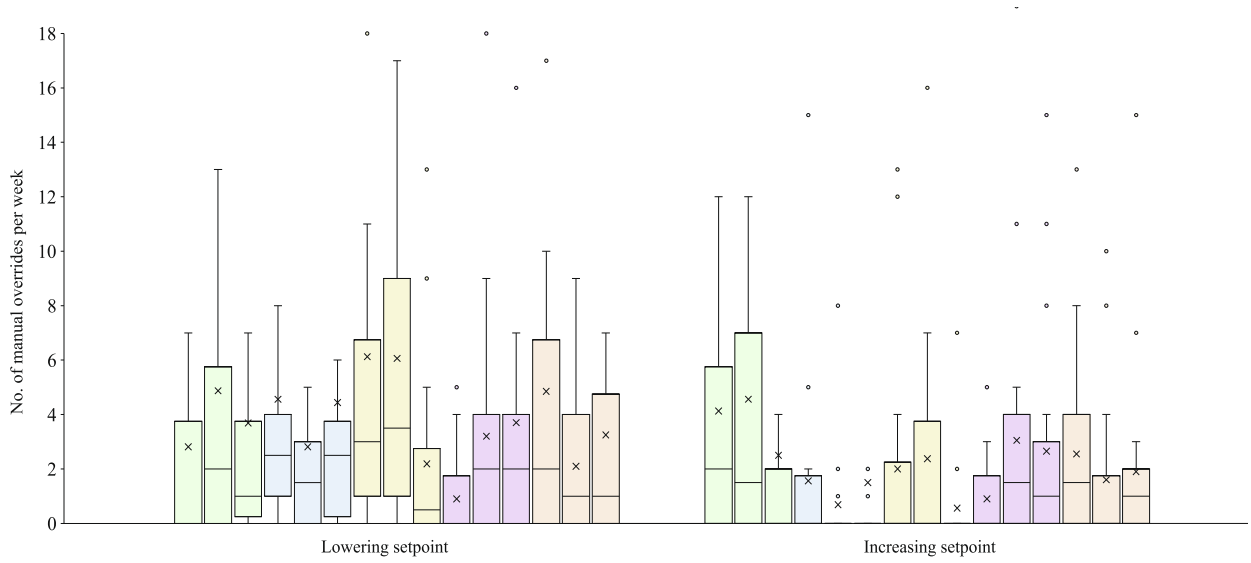

Fig. 15. Manual overrides for all radiator controls during each week of the experiment.

of solar gains on the predominantly south-facing rooms. It was also seen that when given manual control of the radiators, occupants interacted with the radiator more frequently by turning it high and low, trying to find a comfortable spot, since the control was not automatically doing it for them.

3.4. Thermal comfort surveys

A thermal comfort survey was also taken each week for the occupants of each room for the first three control strategies in 2019. Figs. 16, 17, and 18 show the results of this survey, which resulted in between 20 and 22 responses from occupants during these periods. Unfortunately, comfort surveys were not deployed with the PID control strategies implemented due to time and personnel constraints.

From Fig. 16, it is seen that occupants reported the highest level of satisfaction with thermal comfort at the time of the survey (i.e., mid-day) with the manual control (95%) compared to the enforced schedule (75%) and occupancy-based (81%) control strategies. In other words, the enforced schedule and occupancy-based control strategies led to lower radiator runtime but resulted in the percent of people dissatisfied (PPD) in these spot surveys increasing from only 5% (well within the ASHRAE comfort target of < 10%) to 25% and 19%, respectively (both of which are outside of the ASHRAE comfort target of < 10%). This suggests that while energy is conserved, there was some decline in overall comfort, at least at the time of the survey.

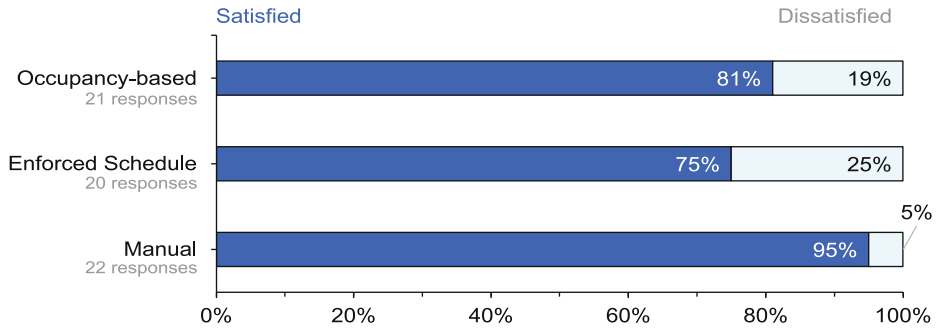

Fig. 16. Thermal comfort survey results from 2019.

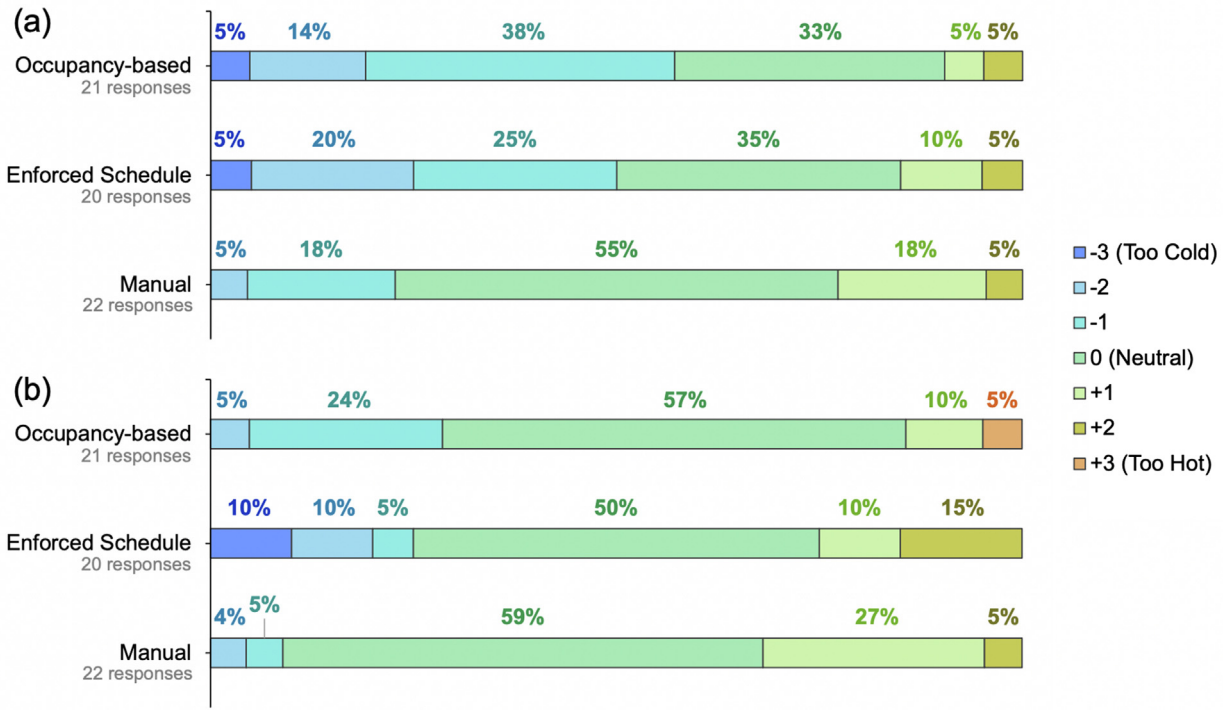

Fig. 17. Thermal comfort sensation analysis from 2019 (a) in the morning; and (b) during the day.

To provide a better understanding of thermal comfort with each control strategy engaged, the occupants were asked how hot or cold they were feeling at the time of the survey (mid-day) and also how they felt that morning when they first arrived in the room. Both questions used a 7-point scale from −3 (too cold) to + 3 (too hot) with 0 indicating neutral or thermal satisfaction. Fig. 17a shows how the occupants felt in the morning and Fig. 17b shows occupants' thermal comfort levels during the time of survey (mid-day).

Occupants mostly reported neutral thermal sensations during the day with all control strategies (50–59%) but reported cooler sensations more often in the morning with the two automatic control strategies compared to the manual control. The mean vote during mid-day periods was 0.34, 0.31, and −0.10 with the manual, enforced schedule, and occupancy-based control strategies deployed, respectively, and 0.27, 0.25, and −0.66 during the mornings of observation days with the manual, enforced schedule, and occupancy-based control strategies deployed. As another measure of perceived comfort, the percentage of occupants reporting thermal sensations between −1 and + 1 during mid-day periods was 91%, 65%, and 92% with the manual, enforced schedule, and occupancy-based control strategies deployed, respectively, and 91%, 70%, and 76% during the same mornings with each strategy deployed.

Investigation of the room level temperature data showed that the warm morning sensations reported by occupants under manual control was largely attributable to the fact that occupants did not turn the radiators off while leaving for the day, so they frequently heated the room throughout the night and into the morning prior to their arrival. This made the room comfortably warm for the occupants at the beginning of the next day, but at the expense of energy being wasted during unoccupied nighttime hours.

The enforced schedule control strategy also caused some occupants to feel slightly warmer than usual on certain days, which can be explained by the variation in their daily schedule as shown previously in Fig. 13. Since the strategy was implemented beginning from 7AM, occupants arriving after 9 or 10 AM usually arrived to a room that was slightly warmer than usual, as compared to those that arrived earlier (e.g., between 7 and 9AM). Those occupants that arrived earlier than most usually felt the room was colder in the mornings. For some occupants, the predetermined setpoint during the enforced schedule control strategy was simply not high

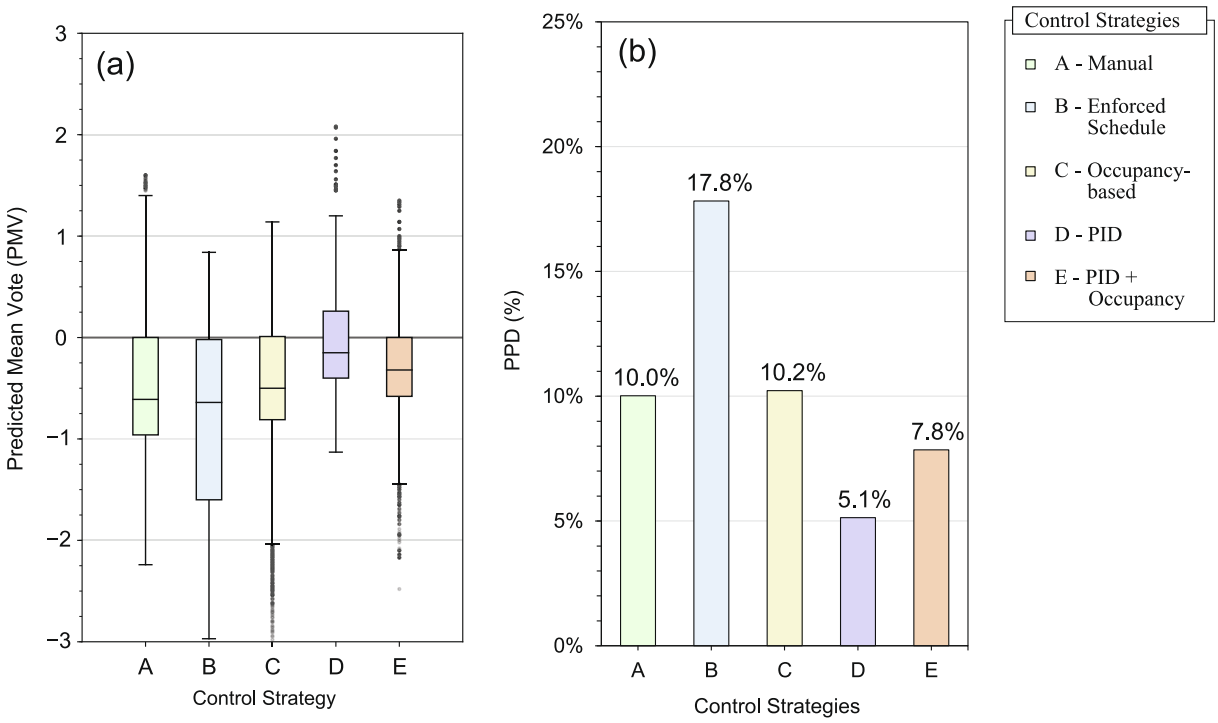

Fig. 18. (a) PMV calculation results; and (b) PPD calculation results.

enough to provide them with sufficient heating during the day. These behaviors can be explained by the fact that the enforced schedule control strategy was generally implemented on much colder days, as shown previously in Fig. 12a.

The occupancy-based control strategy showed the highest percentage of people dissatisfied with cold rooms in the mornings, since the control was activated only when they first entered the room and remained closed when the room was unoccupied. It was only activated when the room temperature was below 22 °C and when the room was occupied (triggered by motion). The combination of runtime results in Fig. 12 and thermal sensation results in Figs. 16-17 suggests that the occupancy-based control could be further optimized to pre-heat rooms prior to arrival while still saving energy throughout non-occupied times. Additionally, while quantitative surveys were not taken for both PID control strategies, regular inquiries with a small number of occupants during each week of the experiment indicated that most occupants were generally very satisfied with the PID control, as supported further with temperature data and predicted comfort in section 3.5.

3.5. Thermal comfort predictions

For a more data-driven thermal comfort analysis, the PMV model was used, informed by measurements of MRT in each room. The PMV index predicts a mean value of subjective ratings of a group of people in a given indoor environment. It is a seven-point scale that represents thermal sensation, ranging from −3 (cold) to + 3 (hot), with 0 representing neutral sensation. PMV indices are used to calculate the Predicted Percentage of Dissatisfied (PPD), which indicates the percentage of people dissatisfied in a given environment. ASHRAE Standard 55–2017 sets the indoor thermal comfort requirements using the PMV model in which at least 80% of the occupants should be satisfied [23].

Using the MRT measurements in each room for each control strategy, PMV indexes were calculated. A Python script was written that uses *pythermalcomfort*, a Python package developed by UC Berkeley [55] that calculates thermal comfort indices using various international standards such as ASHRAE 55–2017 [23] and ISO 7730:2005 [56]. The metabolic rate for all individuals was assumed to be 1.1 Met (~64 W/m^2), typical for people sitting and typing at their desks. The mean clothing insulation factor for the occupants was estimated to be 0.90 clo, slightly lower than the suggested international standard value of 1.0 clo during winters. The clothing insulation was estimated based on observation of occupants and informed surveys. Fig. 18a shows the results of PMV calculations. To predict the percentage of occupants that are dissatisfied with each control strategy, the PPD index was calculated. The PPD calculation results are shown in Fig. 18b.

From Fig. 18a and 18b, it is seen that the manual and occupancy-based control strategies had very similar thermal comfort predictions, with a mean PMV-index of around −0.55. This resulted in an estimated ~ 10% people dissatisfied, which satisfies ASHRAE standard 55–2017 for optimal thermal comfort and is within 10% of PPD reported by actual occupants in Fig. 16. The enforced-schedule control strategy resulted in the lowest thermal comfort (predicted PPD of ~ 18%, which is reasonably consistent with reported PPD in Fig. 16), mostly attributable to a combination of the setpoint not being high enough during the daily heating schedule and the outdoor temperature being unusually low during the winter season.

The PID control had a PMV-index of −0.22, very close to the thermally neutral condition of 0. Although comfort surveys were not deployed during this period, this resulted in around 5% predicted PPD, which is the lowest the scale can go. While Fig. 18a shows that there is a spread of outliers during certain periods, almost all occupants would generally be satisfied by the PID control strategy based on these data and predictions. Adding occupancy-based control to the PID control resulted in a slight increase in the PPD-index and is still well within compliance of ASHRAE Standard 55–2017. The periods of time where occupants were predicted to be uncomfortable were generally during the first few minutes of them entering the room after several hours of it

being unoccupied, which resulted in the radiator turning off the control, thereby cooling the room. As soon as the occupants entered the room, the motion sensor detected their presence, setting the radiator control back to their preferred setpoint which heated the room within minutes. These predictions show that controlling the temperature around the workspace as opposed to controlling the radiator valve itself would lead to optimal thermal comfort for most occupants.

3.6. Estimating radiator energy output

Data from the heat flux sensors were used to analyze and estimate the heat output from the radiators. A sample data set from the heat flux sensors is shown in Fig. 19, where the radiator was turned on for approximately one hour and then turned off till it cooled down to room temperature. The channels indicate each sensor, where Channel 1 represents the sensor placed closest to the radiator valve inlet (as seen on the right in Fig. 8a), and Channel 4 is on the furthest end of the radiator (on the left).

From Fig. 19, we see that when the radiator valve is opened from a cold start, it takes about 5 min until the steam travels all the way through the radiator, and about 20 min for the entire radiator to be uniformly heated. The heat flux output from each segment of the radiator also varies, with the ones closest to the inlet producing the highest heat output, and those furthest away producing the lowest. This could be attributed to heat losses into the room along the length of the radiator, as well as from the back end of the radiator, which faces an uninsulated brick wall from which heat can easily escape. It is also observed that when the valve is closed, the radiator is mostly cooled down evenly, with an average cooldown to room temperature taking about 1.5 h.

To further study the variation in heating time between the various channels, an infrared time lapse of the radiator was taken using a FLIR One Gen 2. This is a miniature infrared camera that attaches to a mobile phone via the USB port. While the accuracy and resolution of this infrared camera is relatively low, the images produced give an adequate understanding of the heat flow through the radiator. Fig. 20 a-e show the sequence of heat flowing through the radiator at 5-minutes interval between each image.

From Fig. 20, we can see that the steam flows from the inlet (Channel 1) horizontally through to the furthest end (Channel 4) of the radiator first, followed by the central regions (Channel 2 and finally Channel 3). This is due to the large pipe that travels the length of the radiator on the top through which steam flows the fastest, and then slowly flowing down each fin of the radiator over time. While the time lag noticed here is significant, most occupants were seated closer to the inlet of the radiator, allowing for faster response time in automated control strategies, especially when using PID control, in which the valve is automatically adjusted frequently to maintain an optimum room temperature.

From the heat flux data and surface area, we approximated the energy output from each radiator during each operational condition. Each radiator in the building had slightly different dimensions based on the size of the room in which it was located, but the radiator sizes, and consequently the number of fins, were largely within 15–20% of each other. Moreover, all radiators were made of the same material (cast iron). Because of this, a simple single radiator model was created, with an estimated average surface area of approximately 5.7 m^2. The instantaneous heat transfer from the radiator to the air inside the room is then given by Eq. (2):

$$Q_{rad} = UA(T_{sur} - T_{air}) \tag{2}$$

where Q_{rad} is the heat output from the radiator (W), U is the thermal transmittance of the radiator surface (assumed constant as 11.3 W/m^2K for cast iron), A is the surface area of the radiator (m^2), T_{sur} is the surface temperature of the radiator (°C), and T_{air} is the temperature of the air in the room (°C). Using this approach, we estimated the heat output from the measured radiators during the time periods in which each control strategy was in place. We estimated that the average radiator heat output was approximately 1560 W between the 19 radiators while operating, varying from under 200 W to a peak of 4200 W at any given time during the entire heating season. This variability accounts for all periods of operation across all radiators.

Next, we calculated total energy output during the time periods in which each control strategy was in place as follows. We calculated the average heat output of all 19 radiators measured in the study at 15-minute time intervals using Eq. (2), which provided a total wattage output value for each interval. We then multiplied that heat output value by 0.25 h to yield heat energy output at each 15-minute interval in units of Watt-hours. The total energy output of the radiators for each control strategy is then given by Eq. (3):

$$E_{cs} = \sum_{t=0}^{n} Q_t \Delta t \tag{3}$$

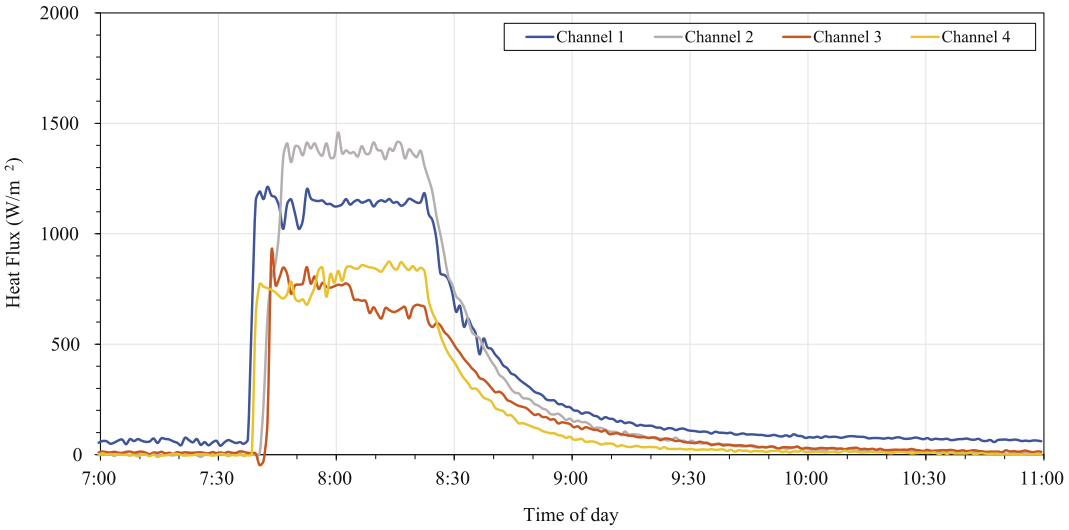

Fig. 19. Sample data from heat flux sensors.

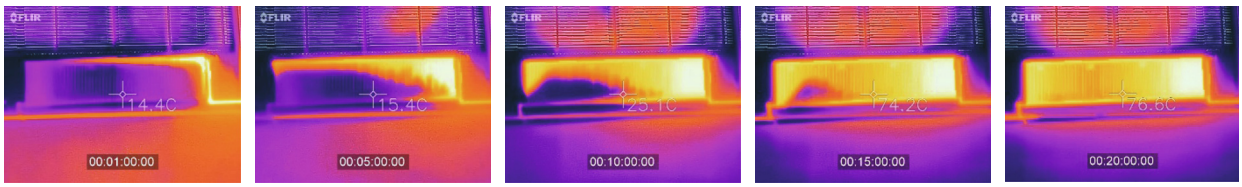

Fig. 20. (a-e) infrared images at 5-minute intervals showing sequence of heat flowing through radiator fins.

where E_{cs} is the energy output of the radiator for each control strategy, Q_t is the average heat output at each 15-minute time interval t (W), Δt is the time interval (15 min or 0.25 h), and n is the total duration of the time period with each control strategy (3 weeks). Table 1 lists the calculated energy use of the radiators for each control strategy for the 3 weeks periods of operation that were measured. Additionally, since the building has 65 radiators in total, with 55 of those being functional and used at least periodically, we also estimated full-building steam usage by scaling the results of our heat output calculations to the whole building for each control strategy, assuming the worst-case scenario of all 55 radiators in use at the same time. These simulated results are a very rough approximation of a worst-case scenario, and while it does not yield reliable quantitative data, it can be used to make qualitative assessments. Therefore, Table 1 does not control for varying temperatures and other factors, but instead utilizes actual measurements of runtime to estimate energy output during each of the 3-week-long test periods.

Using measured runtime and temperature data and estimates of heat output from the 19 actual radiators during the 3-week-long deployment periods, we estimate that the use of the occupancy-based, PID, and PID + Occupancy control strategies would result in energy output savings of approximately 66%, 44%, and 74%, respectively over manual control, while the enforced schedule would not reduce output. Using the cost of steam in Chicago for the year 2019 and 2020 (approximately \$0.027/kWh or \$0.78/therm) [57], we estimate a potential monetary savings of \$452, \$304, and \$514 for the occupancy-based, PID, and PID + Occupancy control strategies, respectively, compared to manual control over the 3-week-long measurement campaigns. For comparison, the total parts cost of outfitting a single zone is approximately \$105, which includes approximately \$70 per fully assembled radiator controller, \$15 for temperature sensing, and \$20 for motion sensing. Thus, the cost of 19 radiator control nodes and temperature and motion monitoring nodes, as well as one gateway, is approximately \$2,000. This estimate does not include labor or time, which required approximately 2 h to assemble and install control nodes, an hour to assemble and program monitoring nodes, and about a day to 3D print parts (which can be automated in batches). Factoring in these costs suggests the controller installation could pay for itself within one or two heating seasons, depending on several factors.

Assuming the worst-case scenario of all 55 radiators in use at the same time, we estimate the impacts of each control strategy on full-building steam usage during the 3-week-long measurement periods to yield monetary savings of up to \$1,310, \$881, and \$1,489 during the 3-week measurement periods for the occupancy-based, PID, and PID + Occupancy control strategies, respectively, compared to manual control. Although these calculations do not include the cost of steam generation, transportation, fuel, boiler efficiency, and energy losses in valves, pressure stations, piping, leaks, and condensation traps, nor do they extrapolate beyond the 3-week operational period of each control strategy or control for weather or occupancy conditions, these estimates do show that adding automation to these legacy systems can yield potential energy savings. The actual energy and cost impacts would be even higher when accounting for all these variables.

4. Discussion

From the measurements and analysis described herein it is seen that the custom radiator controls deployed in this historic building can serve as a viable option to conserve energy use in its legacy steam radiator heating system while also maintaining thermal comfort. These results indicate that such solutions can be used to retrofit existing buildings to convert them to smarter, more efficient buildings, thereby decreasing energy costs without compromising occupant thermal comfort and building system performance. However, this work is not without limitations.

For one, there are many design changes and considerations needed to convert a custom radiator control as described in this paper into a commercial product that can be mass produced. Developing a custom radiator control prototype using 3D-printed parts and other custom-built parts has its compromises. The 3D printed parts were mostly able to withstand the heat from the radiators as well as mechanical stresses from the servo. However, some servos failed due to manufacturing defects and had to be replaced. The failing servos also tend to overheat before failure, which warps the plastic enclosure of the control. The OLED screen used was inexpensive and showed signs of fading on some controls after 2 years of deployment. Some parts of the PCB showed signs of corrosion from being in close proximity to the steam radiators. Moreover, the prototypes described herein require access to electrical

Table 1
Estimated total energy output for radiators for each control strategy during the 3-week-long test periods.

No. of radiators	Measurement	Control Strategy				
		Manual	Enforced Schedule	Occupancy-based	PID	PID + Occupancy
19 radiators (actual)	Total energy output (kWh)	26,035	26,709	8,975	14,562	6,649
	Cost per kWh (\$)	\$0.027	\$0.027	\$0.027	\$0.027	\$0.027
	Estimated cost (\$)	\$690	\$708	\$238	\$386	\$176
55 radiators (simulated)	Total energy output (kWh)	75,365	77,316	25,980	42,152	19,247
	Estimated cost (\$)	\$1,999	\$2,050	\$689	\$1,118	\$510

power, which is not always accessible in buildings with these types of legacy heating systems. Thus, while the prototype devices show promise, we are currently working on an updated version of the controller that can run independently on batteries instead of requiring external power, as well as having self-calibration features and better user interaction.

Additionally, the deployment and evaluation described herein is limited to one building and to only 3-week-long test periods for each control strategy. Further longer-term testing would be useful to quantify potential energy and cost savings and impacts on thermal comfort at a larger scale (i.e., more buildings over longer periods of time). Additional and more detailed simulations of energy use and costs could further strengthen the brief economic analysis provided herein. Moreover, while the PMV/PPD prediction methods used in this study is generally accepted as the industry standard for thermal comfort in buildings, studies have shown that it is not always accurate [58], which creates the need for developing better thermal comfort models which not only take more factors into consideration but also adapt over time to an individual or to a group of people.

Regardless, this work shows that this custom sensing and control solution has the potential to reduce energy use and maintain thermal comfort by leveraging granular zone-level data with wireless sensors, which also provides valuable information about the existing system performance. Such a control system that allows granular setpoint control based on sensor feedback can also be implemented on other systems that incorporate a valve for controlling the flow of fluids, such as floor heating manifolds and other water-based heating systems. This can be done by developing customized mounting hardware for the valves, and control strategies can then be set based on the specifications of the room, along with direct feedback from wireless sensors. In addition to heating system control, this information can be used to identify faulty operation, check when the building systems need tuning, fix operational problems, and reduce inefficiencies by allowing on-demand maintenance and repair beforehand. The proposed control system can also be integrated with forced air HVAC systems to provide on-demand cooling by adjusting supply flow rates and heating by adjusting the radiator valves as shown here for optimal thermal comfort. This can be done by utilizing real-time feedback from HVAC controllers or from BMS, which can be translated into information that the backend can use along with other sensor data to inform the control strategies. Whole building automation can also be achieved for limitless control if the proposed system is allowed to integrate with motorized window shades, lighting, and other mechanical and electrical building systems.

Last, while this work evaluated the impacts of this system during periods of normal occupancy (prior to COVID-19), major changes in occupancy patterns in commercial, educational, and institutional buildings during the COVID-19 pandemic has left numerous buildings unoccupied, which further emphasizes the utility of retrofitting these types of legacy systems with automatic control solutions such as this one.

5. Conclusion

This paper demonstrates an automated radiator control system that can be used to retrofit legacy steam radiator heating systems. The control can be augmented by data from wireless sensors, and it can be optimized for energy savings while also maintaining thermal comfort. Adding automation to these legacy systems can yield potential energy savings and can even provide better thermal comfort when optimized for both endpoints. Most commercial buildings today still operate on an enforced schedule with varying setpoints during the day. This study demonstrates that having a fixed schedule for heating may not always be optimal and may lead to worse thermal comfort in some situations. A combination of real-time occupancy detection and either an on–off controller or a full PID-based controller yields the greatest predicted thermal comfort with the highest predicted energy savings. In future implementations, this setup may be combined with machine learning to create a predictive control that is optimized further over time to constantly strive for that perfect balance between thermal comfort and maximum energy savings as our lifestyle changes and evolves over time.

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.

Acknowledgments

This work was funded by Franklin Energy Services, LLC. Mohammad Heidarinejad was also funded in part by an ASHRAE New Investigator Award. Erica Acton was funded by the Armour College of Engineering's Program for Undergraduate Research in Engineering. We are grateful to Behzad Salimian Rizi for his help in analyzing the steam distribution system in the test building (Alumni Memorial Hall).

Appendix A

A-1 backend server software

To allow easy configuration and control of each control node, a custom backend software solution was developed using WordPress. WordPress REST APIs allow for extremely quick development of custom apps that can interact with any hardware or software. They enable high performance, reliability, and scalability by providing data in a clear manageable format. WordPress provides all accessible data in JSON format, which is universally used by web applications in general. The API used allows for a client–server architecture to be implemented, in which a client initiates a communication session with a server and requests a specific piece of information by accessing an "endpoint". This endpoint is a URL with specific headers to define authentication and type of data requested. The server then responds with data specific to what was requested. WordPress makes developing these endpoints and formatting the actual data extremely easy by using their admin dashboard and allowing customization to it to serve a specific application, in this case, configuring control nodes with custom settings. This eliminates the need to develop a custom web application from scratch and provides tools for developing device specific configuration settings all displayed neatly on a dashboard. An example of such a configuration screen is shown in Fig. A.1.

All configuration settings are stored on the backhaul (where WordPress is installed), and whenever a control node requests its latest configuration, it receives it in a neatly packaged JSON message. An example of such a message is shown below, where each key indicates a specific state, and its corresponding value indicates what the state should be set as.

{
"id": 64,
"acf": {
"control_node_type": "radiator",
"ip_address": "192.168.0.130",

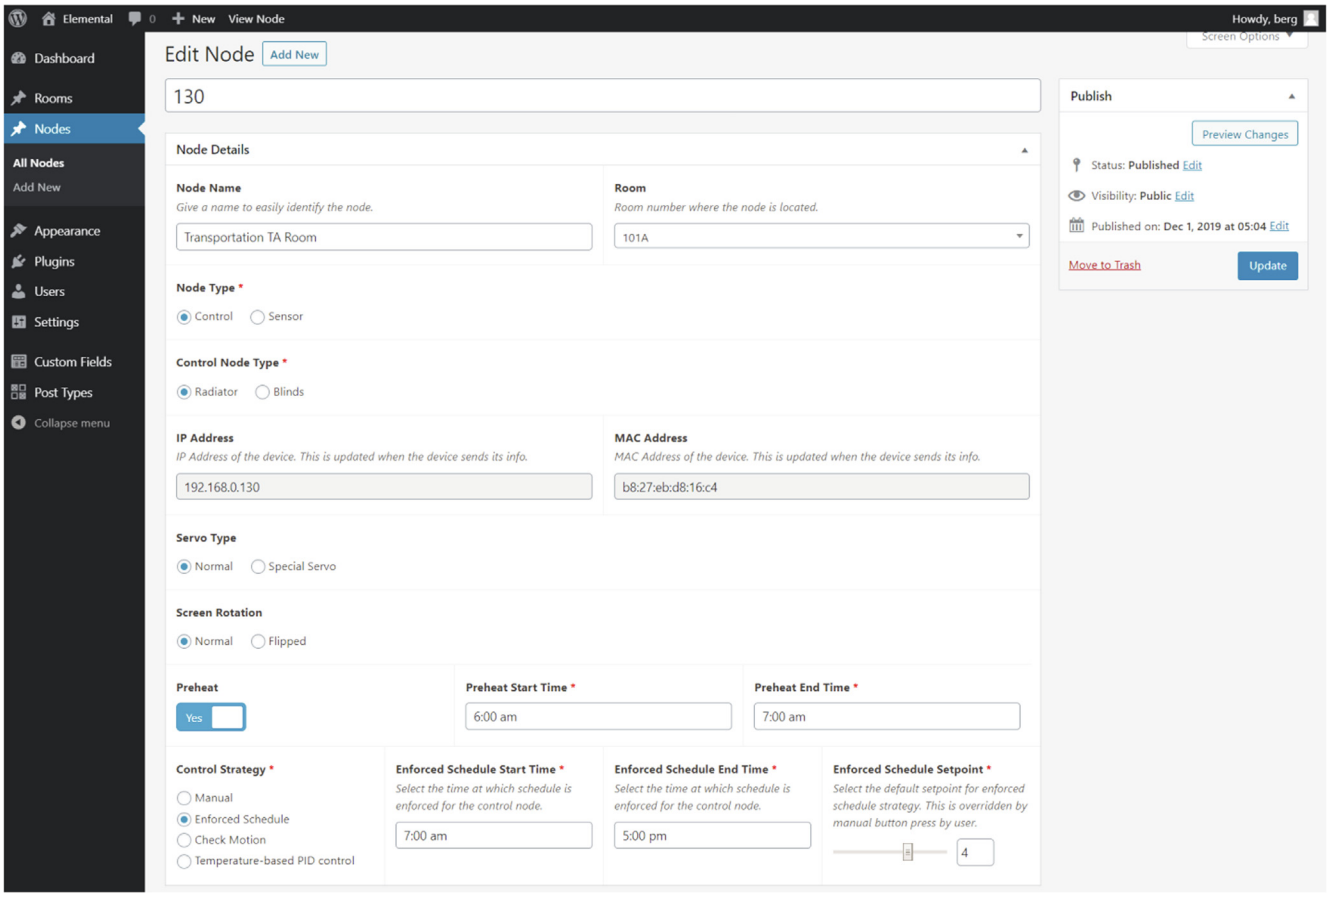

Fig. A1. Web-based configuration screen for a control node.

"mac_address": "b8:27:eb:d8:16:c4",
"servo_type": "normal",
"screen_rotation": "normal",
"preheat": true,
"preheat_start_time": "06:00",
"preheat_end_time": "07:00",
"control_strategy": "enforced_schedule",
"enforced_schedule_start_time": "07:00",
"enforced_schedule_end_time": "17:00",
"enforced_schedule_setpoint": "4",
"check_motion_timeout": "30",
"check_motion_setpoint": "4",
"check_motion_min_temp": "72",
"temperature_setpoint": "72"
}
}

The incoming data are accessible by all Python scripts in the form of a central dictionary saved locally on each control node. This dictionary is updated with new information when requested via the API. Each control node can be individually accessed and customized from the web application to have any desired setting and control strategy, reducing the need to deal with changing code on each device manually. WordPress allows for new control settings to be added and customized on the fly, and this custom control dashboard for the nodes is accessible from any web browser on any device locally or over the internet. This server software can be implemented in the cloud, or locally on any computer dedicated for building controls, or even directly on the Raspberry Pi used as the backhaul, depending on the level of remote access required.

References

[1] N.E.P.Fernandez, S. Katipamula, W. Wang, Y. Xie, M. Zhao, C.D. Corbin, Impacts of commercial building controls on energy savings and peak load reduction; Pacific Northwest National Lab. (PNNL), Richland, WA (United States), 2017.

[2] A. Orestis, A. Dimitrios, D. Dimitrios, C. Ioannis, Smart energy monitoring and management in large multi-office building environments. In Proceedings of the Proceedings of the 17th Panhellenic Conference on Informatics; ACM: New York, NY, USA, 2013; pp. 219–226.

[3] C. de Farias, H. Soares, L. Pirmez, F. Delicato, I. Santos, L.F. Carmo, J. de Souza, A. Zomaya, M. Dohler, A control and decision system for smart buildings using wireless sensor and actuator networks, Trans. Emerg. Telecommun. Technol. 2014, 25, 120–135, 10.1002/ett.2791.

[4] E. Png, S. Srinivasan, K. Bekiroglu, J. Chaoyang, R. Su, K. Poolla, An internet of things upgrade for smart and scalable heating, ventilation and air-conditioning control in commercial buildings, Appl. Energy 239 (2019) 408–424, https://doi.org/10.1016/j.apenergy.2019.01.229.

[5] S. Ferdoush, X. Li, Wireless sensor network system design using raspberry pi and arduino for environmental monitoring applications, Procedia Comput. Sci. 34 (2014) 103–110, https://doi.org/10.1016/j.procs.2014.07.059.

[6] M. Karami, G.V. McMorrow, L. Wang, Continuous monitoring of indoor environmental quality using an arduino-based data acquisition system, J. Build. Eng. 19 (2018) 412–419, https://doi.org/10.1016/j.jobe.2018.05.014.

[7] O. Bamodu, L. Xia, L. Tang, An indoor environment monitoring system using low-cost sensor network, Energy Procedia 141 (2017) 660–666, https://doi.org/10.1016/j.egypro.2017.11.089.

[8] A.S. Ali, Z. Zanzinger, D. Debose, B. Stephens, Open source building science sensors (OSBSS): a low-cost arduino-based platform for long-term indoor environmental data collection, Build. Environ. 100 (2016) 114–126, https://doi.org/10.1016/j.buildenv.2016.02.010.

[9] L. Pocero, D. Amaxilatis, G. Mylonas, I. Chatzigiannakis, Open source IoT meter devices for smart and energy-efficient school buildings, HardwareX 1 (2017) 54–67, https://doi.org/10.1016/j.ohx.2017.02.002.

[10] M.F. AbdRabo, K.M.K. Pasha, Controlling the heat transfer and pressure drop within economical working conditions for a movable flat tube bundle, Int. J. Therm. Sci. 107 (2016) 259–271, https://doi.org/10.1016/j.ijthermalsci.2016.03.022.

[11] J.Y. Park, T. Dougherty, H. Fritz, Z. Nagy, LightLearn: an adaptive and occupant centered controller for lighting based on reinforcement learning, Build. Environ. 147 (2019) 397–414, https://doi.org/10.1016/j.buildenv.2018.10.028.
[12] M.F. Touchie, J.A. Siegel, Residential HVAC runtime from smart thermostats: characterization, comparison, and impacts, Indoor Air 28 (6) (2018) 905–915, https://doi.org/10.1111/ina.2018.28.issue-610.1111/ina.12496.
[13] A. Meier, T. Ueno, M. Pritoni, Using data from connected thermostats to track large power outages in the united states, Appl. Energy 256 (2019), https://doi.org/10.1016/j.apenergy.2019.113940 113940.
[14] A. Elshafee, K. Alaa Hamed, Design and Implementation of a WiFi Based Home Automation System, World Acad. Sci Eng. Technol. 2177 (2012).
[15] Residential Energy Consumption Survey (RECS) - Energy Information Administration Available online: https://www.eia.gov/consumption/residential/ (accessed on 16 December 2020).
[16] D. Schäuble, A. Marian, L. Cremonese, Conditions for a cost-effective application of smart thermostat systems in residential buildings, Appl. Energy 262 (2020), https://doi.org/10.1016/j.apenergy.2020.114526 114526.
[17] Z. Pang, Y. Chen, J. Zhang, Z. O'Neill, H. Cheng, B. Dong, How much HVAC energy could be saved from the occupant-centric smart home thermostat: a nationwide simulation study, Appl. Energy 116251 (2020), https://doi.org/10.1016/j.apenergy.2020.116251.
[18] T. Peffer, M. Pritoni, A. Meier, C. Aragon, D. Perry, How people use thermostats in homes: a review, Build. Environ. 46 (2011) 2529–2541, https://doi.org/10.1016/j.buildenv.2011.06.002.
[19] S. Abbas, A. Bakar, Y. Chandio, K. Hafeez, A. Ali. T.M. Jadoon, M.H. Alizai, Inverted HVAC: greenifying older buildings, one room at a time, ACM Trans. Sens. Netw. 2018, 14, 26:1-26:26, 10.1145/3229063.
[20] N. Li, G. Calis, B. Becerik-Gerber, Measuring and monitoring occupancy with an RFID based system for demand-driven HVAC operations, Autom. Constr. 24 (2012) 89–99, https://doi.org/10.1016/j.autcon.2012.02.013.
[21] L. Klein, J. Kwak, G. Kavulya, F. Jazizadeh, B. Becerik-Gerber, P. Varakantham, M. Tambe, Coordinating occupant behavior for building energy and comfort management using multi-agent systems, Autom. Constr. 22 (2012) 525–536, https://doi.org/10.1016/j.autcon.2011.11.012.
[22] F. Nicol, M. Humphreys, New standards for comfort and energy use in buildings, Build. Res. Inf. - Build. RES Inf. 37 (2009) 68–73, https://doi.org/10.1080/09613210802611041.
[23] ASHRAE Standard 55 – thermal environmental conditions for human occupancy 2017.
[24] S. Baldi, C.D. Korkas, M. Lv, E.B. Kosmatopoulos, Automating occupant-building interaction via smart zoning of thermostatic loads: a switched self-tuning approach, Appl. Energy 231 (2018) 1246–1258, https://doi.org/10.1016/j.apenergy.2018.09.188.
[25] Research and Markets, Ltd. IoT in Smart Buildings Market Outlook and Forecasts 2018 - 2023; Mind Commerce; Global, 2018; p. 274.
[26] T.M. Lawrence, M.-C. Boudreau, L. Helsen, G. Henze, J. Mohammadpour, D. Noonan, D. Patteeuw, S. Pless, R.T. Ten Watson, Questions concerning integrating smart buildings into the smart grid, Build. Environ. 108 (2016) 273–283, https://doi.org/10.1016/j.buildenv.2016.08.022.
[27] American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) 2020 ASHRAE Handbook—HVAC Systems and Equipment; SI Edition.; American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. (ASHRAE): Atlanta, GA, 2020; ISBN 978-1-947192-53-9.
[28] B. Salimian Rizi, A. Syed Ali, C. Riley, B. Stephens, M. Heidarinejad, Energy analysis of steam distribution system using a physics-based model: a campus building case study. in proceedings of the systems, components and loads analysis; ASHRAE and IBPSA-USA: Chicago, IL USA, August 12 2020; Vol. Conference Paper Session 12 (Basic), pp. 463–470.
[29] Demystifying Steam Report (2019) - Urban Green Council Available online: https://www.urbangreencouncil.org/content/projects/demystifying-steam-report (accessed on 16 December 2020).
[30] J. Michaels, Energy Information Administration (EIA)- Commercial Buildings Energy Consumption Survey (CBECS) Available online: https://www.eia.gov/consumption/commercial/ (accessed on 6 February 2020).
[31] P. Lauenburg, J. Wollerstrand, Adaptive Control of radiator systems for a lowest possible district heating return temperature, Energy Build. 72 (2014) 132–140, https://doi.org/10.1016/j.enbuild.2013.12.011.
[32] R. Ruch, P. Ludwig, T. Maurer, Balancing hydronic systems in multifamily buildings; 2014; p. DOE/GO–102014-4463, 1221046.
[33] M. Bobker, E.L. Kinsley, Balancing apartment building heating with thermostatic radiator valves, Heat. Piping Air Cond. 67 (1995).
[34] Thermostatic Radiator Valve (TRV) Demonstration Project. Final Report; New York State Energy Research and Development Authority, Albany, NY (United States); EME Group, New York, NY (United States), 1995.
[35] T. Cholewa, A. Siuta-Olcha, C.A. Balaras, Actual energy savings from the use of thermostatic radiator valves in residential buildings – long term field evaluation, Energy Build. 151 (2017) 487–493, https://doi.org/10.1016/j.enbuild.2017.06.070.
[36] B.-C. Ahn, J.-Y. Song, Control characteristics and heating performance analysis of automatic thermostatic valves for radiant slab heating system in residential apartments, Energy 35 (2010) 1615–1624, https://doi.org/10.1016/j.energy.2009.11.007.
[37] J.Y. Park, Z. Nagy, Comprehensive analysis of the relationship between thermal comfort and building control research - a data-driven literature review, Renew. Sustain. Energy Rev. 82 (2018) 2664–2679, https://doi.org/10.1016/j.rser.2017.09.102.
[38] B. Xu, L. Fu, H. Di, Dynamic Simulation of space heating systems with radiators controlled by TRVs in buildings, Energy Build. 40 (2008) 1755–1764, https://doi.org/10.1016/j.enbuild.2008.03.004.
[39] V. Monetti, E. Fabrizio, M. Filippi, Impact of low investment strategies for space heating control: application of thermostatic radiators valves to an old residential building, Energy Build. 95 (2015) 202–210, https://doi.org/10.1016/j.enbuild.2015.01.001.
[40] K.K. Andersen, H. Madsen, O. Knop, N. Gregersen, Modelling of a thermostatic valve with hysteresis effects, IFAC Proc. 35 (2002) 173–178, https://doi.org/10.3182/20020721-6-ES-1901.01582.
[41] P. Weker, J.M. Mineur, A performance index for thermostatic radiator valves, Appl. Energy 6 (1980) 203–215, https://doi.org/10.1016/0306-2619(80)90034-3.
[42] A. Bruce-Konuah, R.V. Jones, A. Fuertes, L. Messi, A. Giretti, The role of thermostatic radiator valves for the control of space heating in UK social-rented households, Energy Build. 173 (2018) 206–220, https://doi.org/10.1016/j.enbuild.2018.05.023.
[43] M.K. Fahnestock, Effect of Enclosures on Direct Steam Radiator Performance. Univ. Ill. Urbana Champaign Coll. Eng. Eng. Exp. Stn. 1927.
[44] Radiator Labs, Inc. | Smart Radiator Cover, The Cozy.
[45] T. Manning, R.T. Nortcliffe, M.I. Wilson, S.J. Mitchell, Thermostatic radiator valves and control thereof 2010.
[46] H.A. Thiel, Control device for radiator valves 2005.
[47] S.H. Hong, S.H. Kim, G.M. Kim, H.L. Kim, Experimental evaluation of BZ-GW (BACnet-ZigBee Smart Grid Gateway) for demand response in buildings, Energy 65 (2014) 62–70, https://doi.org/10.1016/j.energy.2013.12.008.
[48] A.S. Ali, C. Coté, M. Heidarinejad, B. Stephens, Elemental: an open-source wireless hardware and software platform for building energy and indoor environmental monitoring and control, Sensors 19 (2019) 4017, https://doi.org/10.3390/s19184017.
[49] ASHRAE 102-1983 (RA 89) 1983.
[50] E.W. Shaw, Thermal comfort: analysis and applications in environmental engineering, by P. O. Fanger. 244 Pp. DANISH TECHNICAL PRESS. Copenhagen, Denmark, 1970. Danish Kr. 76, R. Soc. Health J. 1972, 92, 164–164, 10.1177/146642407209200337.
[51] 14:00-17:00 ISO 7726:1998 Available online: https://www.iso.org/cms/render/live/en/sites/isoorg/contents/data/standard/01/45/14562.html (accessed on 20 August 2020).
[52] PID Controller. Wikipedia 2019.
[53] Y. Yau, B.T. Chew, A.Z.A. Saifullah, A field study on thermal comfort of occupants and acceptable neutral temperature at the national museum in malaysia, Indoor Built. Environ. 22 (2013) 433–444, https://doi.org/10.1177/1420326X11429976.
[54] Local Weather Forecast, News and Conditions | Weather Underground Available online: https://www.wunderground.com/ (accessed on 21 October 2020).
[55] F. Tartarini, S. Schiavon, Pythermalcomfort: a python package for thermal comfort research, SoftwareX 12 (2020), https://doi.org/10.1016/j.softx.2020.100578 100578.
[56] 14:00-17:00 ISO 7730:2005 Available online: https://www.iso.org/cms/render/live/en/sites/isoorg/contents/data/standard/03/91/39155.html (accessed on 20 August 2020).
[57] Natural Gas Data - U.S. Energy Information Administration (EIA) Available online: https://www.eia.gov/naturalgas/data.php (accessed on 3 January 2021).
[58] T. Cheung, S. Schiavon, T. Parkinson, P. Li, G. Brager, Analysis of the accuracy on PMV – PPD model using the ASHRAE global thermal comfort database II, Build. Environ. 153 (2019) 205–217, https://doi.org/10.1016/j.buildenv.2019.01.055.