# OcAPO: Occupancy-aware, PDC control for open-plan, shared workspaces

Anaradha RAVI

Archan MISRA
*Singapore Management University*, archanm@smu.edu.sg

# OcAPO: Occupancy-Aware, PDC Control for Open-Plan, Shared Workspaces

Anuradha Ravi
*School of Computing and Information Systems*
*Singapore Management University*
anuradhar@ieee.org

Archan Misra
*School of Computing and Information Systems*
*Singapore Management University*
archanm@smu.edu.sg

***Abstract*—Passive Displacement Cooling (PDC) has gained popularity as a means of significantly reducing building energy consumption overheads, especially in tropical climates. PDC eliminates the use of mechanical fans, instead using chilled-water heat exchangers to perform convective cooling. In this paper, we evaluate the impact of different parameters affecting occupant comfort in a $1000m^2$ open-floor area (consisting of multiple zones) of a ZEB (Zero Energy Building) deployed with PDC units and tackle the problem of setting the temperature setpoint of the PDC units to assure occupant thermal comfort. We tackle two key practical challenges: (a) the zone-level (i.e., occupant-experienced) temperature differs significantly, depending on occupancy levels, from that measured by the ceiling-mounted thermal sensors that drive the PDC control loop, and (b) sparsely deployed sensors are unable to distinguish between ambient temperature variations across neighboring zones. Using extensive real-world measurement data (collected over 60 days), we devise a trace-based model that helps identify the optimum combination of PDC setpoints, collectively across multiple zones, while accommodating variations in the occupancy levels and weather conditions. We deploy *OcAPO* on our real-world testbed to demonstrate its efficacy: while *OcAPO* reliably assures occupancy comfort within a tolerance of $0.2°$C, the current practice of occupancy-agnostic rule-based setpoint control violates this tolerance value 75.2% of the time.**

***Index Terms*—HVAC control, thermal comfort, Smart Building Management, occupancy estimation**

## I. Introduction

Commercial buildings generate $\sim$33% of the world's total energy consumption, within which HVAC represents the dominant (38%, translating to 12% of the global energy demand) consumption load [1]. Passive Displacement Cooling (PDC) technologies[1] have gained significant transaction as an energy-efficient cooling mechanism (reducing the cooling energy overhead by $\sim$15%), especially in tropical regions where heating support is typically unnecessary. PDCs fundamentally dispense with the use of energy-consuming fans for forced air circulation, and instead utilize the principles of natural convective cooling, where chilled air (typically cooled by heat exchange with chilled water pipes) drifts towards the lower portions of the room while warmer air rises towards the ceiling. Combining such PDC with occupancy-aware control [2], can further lower the overall energy cost.

Early research on occupancy-aware HVAC control relied principally on a binary estimator occupancy (occupied vs. unoccupied), often using PIR [3] or $CO_2$ sensing [4] to estimate whether a region has *any* occupants or not, and accordingly activate or deactivate cooling functions. To overcome the natural transients (hysteresis) associated with cooling, more sophisticated methods employ occupancy prediction techniques [5, 6] to control HVAC valves preemptively, initiating the cooling down of a space before the anticipated arrival of human occupants. These strategies have, however, been predominantly tested in traditional *compartmentalized* commercial spaces, where individual layouts (e.g., office or meeting rooms) are effectively thermally isolated.

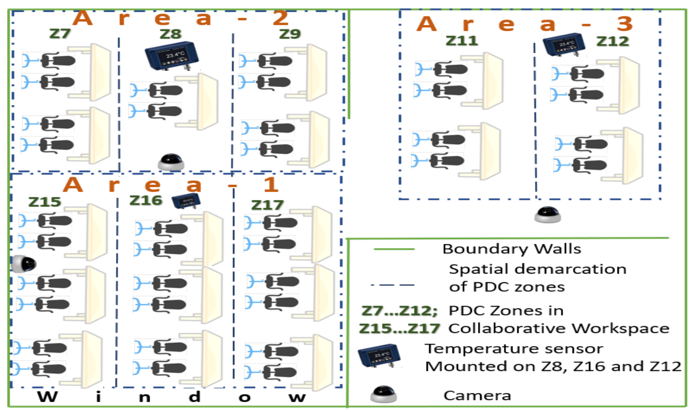

Fig. 1: Layout of Collaborative Workspace (Testbed)

In this work, we tackle the problem of performing efficient, occupancy-aware PDC-based cooling in more modern, open-plan workspaces where the lack of explicit walls/partitions implies a strong mutual correlation between the thermal states of nearby or neighboring spaces. Our goal is to intelligently adjust the temperature set points, which in turn activate or deactivate PDC controllers, so as to assure occupant thermal comfort in such open spaces where the aggregate number of occupants can exhibit significant variation. As a specific exemplar of this problem, we tackle the issue of occupancy-aware, thermal comfort preserving, PDC control of a $\sim 1000m^2$ open-plan collaboration space on one floor of an operational university ZEB (Zero-Energy Building). Figure 1 illustrates the layout of our collaborative workspace, which consists of 8 zones (virtually divided into 3 areas), each individually controlled by a ceiling-mounted PDC air vent that chills cold air that drifts downwards, while the hot air rises. For this specific deployment, the PDC vents open up to circulate the cool air when the *nominal* room temperature, measured via

[1]https://www.daikin.com.sg/product-series/pdv/

ceiling-mounted temperature sensors, exceeds the zone's PDC setpoint and closes when this temperature falls below the zone's PDC setpoint.

The collaboration space presents a couple of difficult challenges that we believe to be representative of legacy environments where the PDC system, including vents and PLC control logic, have been deployed by commercial vendors and which we now seek to augment with more sophisticated occupancy-aware control techniques. (Note that the only means of programmatic external control is via setpoint alteration, as the underlying low-level PLC control logic is not accessible.)

- The temperature sensors, whose readings are compared against the programmed set points, are mounted on the ceiling. Both our empirical studies (detailed in Section IV-B) and prior research [7] provide evidence that the actual floor/table-level ambient temperature, experienced by the human occupants, is higher than the temperature measured at the ceiling. Moreover, we shall show (Section IV-A) that this table-vs.-ceiling discrepancy is a function of different occupancy *levels*, with the ambient temperature being much higher under high occupancy.
- While the 8 zones in the exemplar workspace each have their individual PDC controller, the ceiling-level temperature sensor is deployed more sparsely, with one temperature sensor serving multiple (typically 2-3) zones. Given the open-plan layout of the space, this implies that the reading captured by the sensor is also a *combined* function of the possibly different temperature values in each of the constituent zones.

Any practical solution for dynamic setpoint adjustment must thus be able to (a) accommodate the varying occupancy levels of different zones, (b) estimate each zone's table-level temperature from the ceiling-level thermal sensor readings, and (c) control a *set* of setpoints jointly.

We present our approach, called *OcAPO*[2], that effectively tackles these challenges. Broadly, *OcAPO* operates as follows. First, by utilizing extensive empirical training data, it builds a lookup table that correlates the desired table-level zone-level temperature to appropriate temperature setpoint values under varying levels of $\{$*occupancy levels, external weather conditions*$\}$. Subsequently, during operational deployment, we build a practical 15-20 minute lookahead, occupancy prediction system that combines real-time, non-binary occupancy estimates (derived by executing vision-based DNNs on infrastructural camera feeds) with historical occupancy data to predict zone-level occupancy. Finally, using the occupancy and external weather measurements, *OcAPO* derives and sets temperature setpoints for multiple adjacent zones collectively so as to minimize a measure of *occupant thermal discomfort*.

**Key Contributions:** We make the following key contributions:

- Using extensive real-world measurements and observational data, we quantify both how (a) the relationship between a setpoint value and the occupant-perceived (table-level) ambient temperature is occupancy level dependent, and (b) the occupant-perceived ambient temperature differs (is usually higher than) the ceiling-level temperature sensor readings under varying levels of occupancy. These findings not only motivate the need for an occupancy-aware PDC control logic, but also reveal the need to move beyond simple binary (occupied vs. unoccupied) occupancy estimates.
- We demonstrate the practical challenge of non-binary occupancy prediction (i.e., over 15-20 min lookahead intervals) for different regions of an open plan university collaboration space, where student movement patterns exhibit a high level of spontaneity. To tackle this challenge, we instead combine the diurnal occupancy trend (number of occupants in the morning/afternoon/evening) with real-time video sensing-based occupancy estimation to obtain reasonably accurate predictors for occupancy 15-20 mins in advance.
- We construct a trace-based model to find, given the weather condition and estimated per-zone occupancy, the optimum tuple of zone setpoints that minimize the cumulative deviation from a target ambient temperature (23.5°C). The model utilizes a lookup table, where the occupancy is expressed in 4-levels (unoccupied, low, medium, high) and the weather conditions are represented by 2°C intervals. This approach results in relatively low error across a range of occupancy conditions, achieving mean error of 0.3°C under low, 0.3°C under medium, and 0.1°C under high occupancy levels.
- We conduct pilot studies on the ZEB collaborative space, evaluating the efficacy of *OcAPO's* occupancy-aware dynamic setpoint control logic vs. the baseline of constant 24°C setpoint, over a span of 3 days under natural, real-world occupancy conditions. We empirically show how dynamic occupancy-aware setpoint adaptation is important for assuring occupant comfort: over a total deployment period of 14.5 hours, *OcAPO* specified setpoint readings that are distinct from the ambient objective of 24°C for 70.6% of the time–i.e., a total of 636 minutes and resulted in a mean cumulative deviation of 0.2°C. In contrast, the baseline constant setpoint method exceeded the tolerance range 23.0-24°C 75.2% of the time, with a mean deviation of 0.8°C.

## II. Related Work

The evidence of research in demand-driven HVAC control based on occupancy and weather conditions goes back for decades [8]. However, the evolution in hardware and the subsequent establishment of programmatic interfaces to control HVAC give rise to new challenges and possibilities. Dong et. al. [9] utilized the indoor temperature and CO2 levels, indoor and outdoor RH levels and occupancy state provided by a motion sensor to perform non-linear model predictive controls of HVAC parameters. PIR sensors have proven to reduce energy consumption by 18% [10] in long-term usage. Researchers have also used other sensing modalities, such as WiFi [11] and vision, to estimate occupancy and perform HVAC actuation. Nagarathinam et. al. [12] discuss the possibility of maintaining optimum thermal comfort by modeling the mean radiant temperature of the room at different

[2]**OC**cupancy-aware, PDC Control for **O**pen-Plan, **S**hared Workspaces

occupancy levels, and appropriately adjusting each AHU's (Air Handling Unit) operating parameters. They demonstrate an energy savings of 15%.

Some early research [13], [14] has applied Markov and Regression-based models to control HVAC, taking into account the current and predicted occupancy. Researchers have developed model predictive control (MPC) strategies [6], [15] [16] to perform real-time HVAC adaptation, while adhering to system constraints. Smarra et. al. [17] use random forest-based MPC for HVAC control, achieving an impressive energy savings of 49.5%.

More recently, reinforcement learning (RL) based approaches have been proposed to support *data-driven* HVAC control. Various machine learning algorithms (e.g., [18, 19, 20], [21], [22] [23]) have been propposed to learn the dependencies between various features affecting the thermal comfort of users. Chen et.al. [24] developed an RL-based differential MPC policy for HVAC control that reduced average HVAC energy consumption by 16.7%.

Despite this extensive body of research using state-of-the-art techniques, the actual deployment of extensive sensor instrumentation to support smart HVAC control has been muted. Operational deployments, such as the real-world university space that we shall experimentally study, primarily use motion and temperature sensing input as part of a *rule-based* PDC control policy. In our rule-based deployment, the PDC setpoint is set a static value of 24°C, with the PDC valve closing when no motion is found for 30mins and re-opening whenever the motion is subsequently detected. Motion sensors, however not only often fail to detect stationary humans but also provide only binary occupancy status–this significantly limits the efficacy of the PDC systems. To overcome these shortcomings, we shall adopt a data-driven, non-binary occupancy based PDC control logic.

## III. TestBed and Data Collection

We first describe our target operating space and the data collection mechanism for PDC setpoint control.

### A. TestBed - Collaborative Zone

Our testbed is a collaborative shared workspace of area $1000m^2$ used by students of our university campus. It is logically divided into 8 zones virtually represented as 3 different areas, with no physical demarcation between them (Fig 1). A building management system (BMS) manages each zone equipped with PDC units for cool air circulation at the ceiling level. The PDC relies on the natural convection process of pushing the hot air up towards the ceiling generated by the occupants OR other heat-generating equipment/devices in the space while pushing down the cold air towards the ground level. The cold air is driven by the PDC valves, which are opened between 0-100% based on an internal PLC-based PID (proportional integral derivative) controller. The valve behavior depends on the difference between the ceiling-mounted (room) temperature sensor reading and the setpoint of each zone: the valves open if the room temperature is above the PDC setpoint, and close if the room temperature drops below this setpoint. Because this is a real-time operational building, we are constrained to adjust the PDC setpoint between 22°C and 26°C and do not have any direct control over the opening/closing of PDC valves. In our testbed, PDC valves open up during the operational hours of the building (8.30 AM to 10.30 PM on weekdays). The testbed is characterized by a single ceiling-mounted temperature sensor providing *nominal* temperature readings for each area, comprising multiple zones: Area-1 consisting of Z15, Z16 and Z17 (sensor mounted on Z16), Area-2 consisting of Z7, Z8 and Z9 (sensor mounted on Z8) and Area-3 consisting of Z11 and Z12 (sensor mounted on Z12).

### B. Data Collection

Our testbed is equipped with Beckhoff PLC and is interfaced with our research servers using the OPC-UA interface. We subscribe to the relevant OPC tags (room temperature (RTTS), setpoint value (CRSP), and valve control (VALC)); the PLC notifies our server of updated readings whenever the value of one of these tags changes. The existing infrastructure does not record the fine-grained changes at the zone or occupant level. Hence, we collect fine-grained table-level (ambient) temperature data for each zone by placing two tripod-mounted Texas instrument BLE sensor tags per zone [3]. Each sensor is connected to a Raspberry Pi device and transmits temperature data once every minute via Bluetooth. We then averaged the temperature from 2 sensors of each zone to obtain the zone-level *ambient* temperature data. We collected data intermittently over 4 months between February - March 2022 for measuring the temperature with different occupancy levels and May - June 2022 during the holiday season when there were no occupants.

During this period, we collected the room temperature data to understand (i) the effect of varying weather conditions, (ii) the effect of varying occupancy while adjusting PDC setpoints between $22.5°$C and $25.0°$C in increments of $0.5°$C, and (iii) the effect of PDC setpoint modification on each zone and its neighboring zones. To evaluate the effect of different parameters, we vary the setpoints at different times of the day (Morning, Afternoon, and Evening). While we had the liberty of varying all the PDCs concurrently when the space was unoccupied, we had to be mindful of preserving occupant comfort whenever the space was occupied. Hence, we always kept one zone in each area at $24°$C while varying the PDC setpoint of the other zones in each area between $22.5°$C and $25°$C. For each area, we collected the temperature data for different combinations of zone setpoints and under different occupancy levels, with every zone experiencing a temperature between $22.5 - 25°$C. However, we did not experience high occupancy in the morning and evening times. To understand the phenomena of change in room temperature for varying weather conditions, we also crawled the temperature data for the weather station closest to the academic building

[3]https://www.ti.com/tool/TIDC-CC2650STK-SENSORTAG

from the national real-time weather database [4]. To estimate the occupancy, we use feeds from existing ceiling-mounted cameras to count the number of people in each zone. We use a state-of-the-art object detector (YOLOV3) to compute bounding boxes across persons detected in the scene. We extract the object's pixel coordinates (Xmin, Ymin, Xmax, Ymax, Xcenter, Ycenter) and localize the bounding boxes to the zone's physical coordinates using a logistic regressor.

## IV. Empirical Findings on PDC Setpoint Control

We now discuss several key findings on how occupancy and ceiling-based sensor readings affect occupant thermal comfort.

### A. Impact of Occupancy and Weather Conditions

We individually evaluate the impact of occupancy and weather conditions on the room temperature (measured at both the ceiling and table/occupant levels). Although it is well-established that, in principle, occupancy and weather impact the room temperature, we conduct a fine-grained analysis of how the ambient (occupant-level) temperature is affected by these factors. We first find all the dates and timeslots where the PDC setpoints remained constant, for a given area, throughout the day. We then bin the weather conditions and occupancy values, by (a) quantizing occupancy into 4 distinct levels: {*No*: 0; *Low*:1-3; *Med*:4-7; *High*:>8}, and (b) quantizing external temperature into 6 bins, {22:22.0-24.0; 24:24.1-26.0; 26:26.1-28.0; 28:28.1-30.0; 30:30.1-32.0; 32:32.1-34.0}, which collectively capture the climatic variation of temperature observed. For comparative analysis, we first identify similar weather conditions across the four months (Feb-Mar and May-Jun 2022) under zero occupancy, with the observation period classified into Morning (8-13:00 Hrs), Afternoon (13-18:00 Hrs), and Evening (18-23:00 Hrs).

We track both the zone-level ambient (measured using BLE-equipped TI temperature sensors) and the ceiling-level temperature values. For exposition, we use a representative combined PDC setpoint value of (25, 24, 25) for Area-1 (Z15-16-17). Fig 2 and Fig 3 illustrate the recorded temperature for two different weather buckets, namely, $\sim 25°$C and $\sim 33°$C, with this area completely unoccupied. We observe that the average ceiling temperature recorded at 25°C is $\sim$24.1°C, and the valves were only open for an average duration of 70% (varying between 1%-80%). At 33°C, the ceiling temperature rises up to 24.5°C, and the valves remain open throughout at 100%. While the ambient temperature was significantly different for external temperatures of 24°C and 32°C, the differences were marginal between 22° and 28°C. Fig 4 and Fig 5 illustrate the variation in room temperature for two different occupancy conditions, using the setpoint temperature setting of (25, 24, 25) for the three zones Z15-Z16-Z17. We observe that the average ceiling temperature for low occupancy was recorded to be 25.0°C, while for high occupancy, the ceiling temperature remained at 24.7°C (as valves opened up due to high temperature). We also note that the ambient zone temperature recorded under low occupancy is significantly lower than that at high occupancy. The valves for Z15 and Z17 remained shut throughout the low occupancy, while for high occupancy, the valves opened up to 50% as the ceiling temperature increased up to 25.2°C.

[4] https://data.gov.sg/dataset/realtime-weather-readings

### B. Impact of PDC Setpoint on Ceiling and Zone Temperature

With the objective of understanding the coupled effect of PDC setpoint values across adjacent zones, we simultaneously adjusted the PDC setpoint of two zones (Z15 & Z17 for Area-1, and Z7 & Z8 for Area-2) in decrements of 0.5°C every 30 mins, starting from 25C, and subsequently reversing the steps (in 0.5°C increments) back to 25°C. Because the experiment was performed in an occupied setting, we maintained a consistent temperature of 24C for Z16 in Area-1 and Z9 in Area-2. For Area-3, we varied the temperature for both zones. All the experiments were performed under high occupancy conditions. Fig 6 illustrates the variation in room temperature (ceiling and zone level) for all three areas. We list our observations as follows:

- **Gradual change in zone temperature:** A change in PDC setpoint starts to manifest an impact on occupant-level temperature values after a gap of $\sim$20 mins. However, it can take longer if the setpoint temperature is higher than the room temperature. For instance, as observed in Area-3 (Z11-12), the temperature drops only after the setpoint temperature was set at 24.0°C. This is unlike the phenomena illustrated in Area-1 (Z15-16-17), where the temperature drops after 20 mins immediately after the setpoint temperature is reduced from 25°C to 24.5°C.
- **Impact on Nearby Zones:** For Area-1, we observe that the ambient temperature for zone Z16 reaches the optimum temperature of 24°C only when we lower the PDC setpoints of the other two zones. The zone level temperature of Z15 and Z17 does not reach the optimum temperature in this experiment. However, we observe that the temperature reaches 23.5°-24°C if the zone's setpoint is retained at a lower temperature for a longer duration. For Area-3, we observe that occupant level temperature drops to the optimum level of 24.0°C when both zones are set at 23°C. The data reveals the non-trivial effects of thermal coupling between zones, which motivates us to adopt a strategy of adjusting zone setpoints *collectively*, rather than individually.
- **Varying relation between ceiling and zone temperature:** The ceiling temperature for Area-2 (Z7-8-9), and Area-3 (Z11-12) is significantly lower than the occupant-level temperature. We hypothesize that this phenomena, observed consistently across multiple areas, is because (a) the density of occupants in the two areas is, on average lower by 70% compared to Area-1, and (b) the ceiling temperature sensor is mounted on one of the unoccupied zones of both areas, and thus unable to accurately capture the heat generated within the occupied zone. As developing a precise mathematical model relating ambient to ceiling-level zone temperature seems very challenging and would require significantly more data, we instead design *OcAPO* to adopt a more empirical data-driven approach.

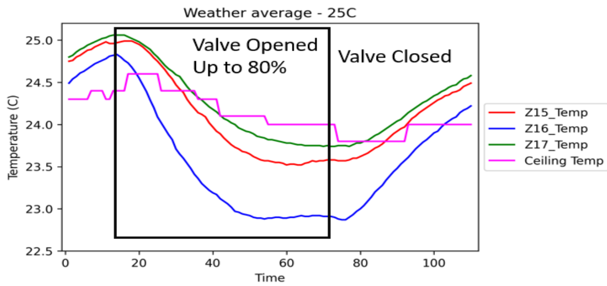

Fig. 2: Weather Average Temperature - 25C

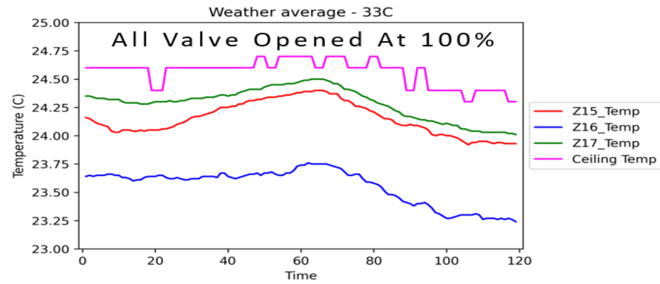

Fig. 3: Weather Average Temperature - 33C

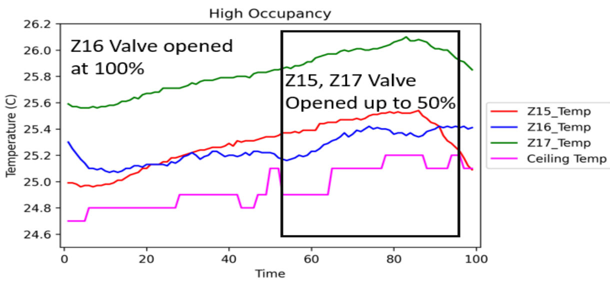

Fig. 4: Room Temperature during high occupancy

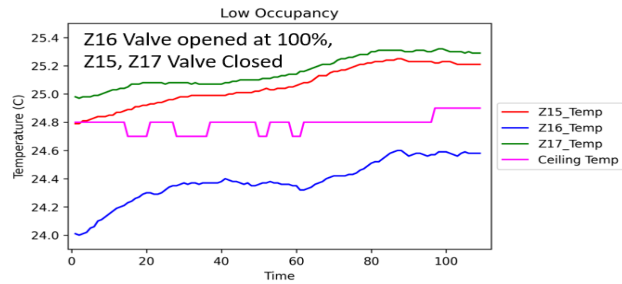

Fig. 5: Room Temperature during low occupancy

Fig. 6: Recorded Room Temperature for Varying PDC Setpoints

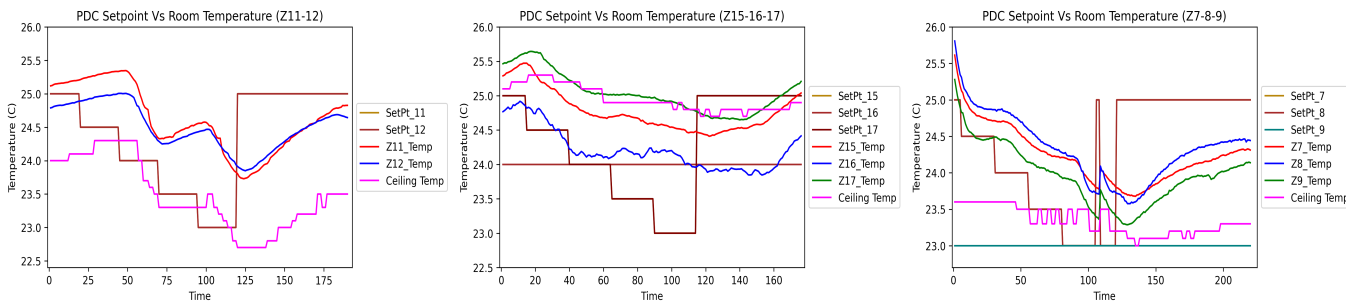


## V. Proposed PDC Setpoint Control

Ideally, PDC setpoint control should use a time series model that factors in the time-varying occupancy levels and the gradual evolution in the room temperature data. However, building such a supervised time series model requires a significant amount of data, capturing fine-grained variations in ambient room temperature, external weather and occupancy. For practical use, we instead develop a trace-based model that uses the observed testbed data to build a lookup table that captures the first-order relationship between these attributes. We then use diurnal models of occupancy and real-time occupancy estimates, obtained by applying object detection vision DNNs over camera data (images captured via CCTV cameras in the infrastructure), together with a *comfort cost function* (CCF), to determine the optimal tuple of setpoints over small (15 minute) time intervals.

### A. Short-Term Occupancy Prediction

As per our analysis reported in Section IV-B, it takes ∼20 mins to cool an area. Accordingly, our approach chunks time in 20 min windows (or epochs), and seeks to perform one-epoch lookahead occupancy prediction (i.e., 20 mins ahead of time). However, predicting such coarse-grained zone-level occupancy using historical data at 15-20 minute temporal granularity is challenging due to unpredictable fluctuations both in (a) the number of occupants in each zone, (b) their arrival time and stay duration. However, we do observe a simple diurnal trend for each zone across different day segments–morning (8am-1pm), afternoon (1pm-3.30; 3.30-6.30) and evening (6.30-10.30pm). For example, Figure 7 plots the diurnal occupancy over one week for Area-1 (comprising zones Z15, Z16 and Z17). In addition, for our analysis and model building, we use the 4 distinct occupancy levels {No, Low, Medium and High} explained earlier. We see that while no clear short-term trend (defined by the divergence from a moving average value) is visible, the occupancy typically is Low-Medium in the morning, High in the afternoon and Low in the evening. Accordingly, at the start of daily segment, we use the observed diurnal value; for all subsequent 20-min intervals in that segment, we approximate the predicted one-step lookahead occupancy based on the current epoch's real-time occupancy estimate (obtained from the camera feeds).

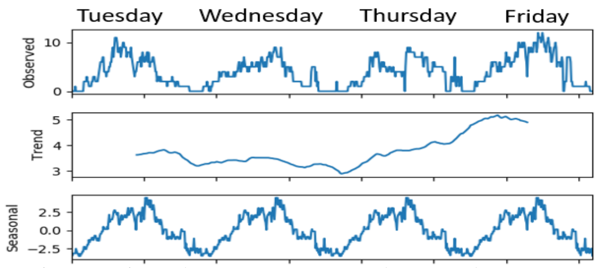

Fig. 7: Diurnal Occupancy Trend (1 week, Area 1)

### B. Lookup Table For Trace-Based Model

Using the data collected over 60 days, we populate a lookup table that helps compute the optimum combination of PDC setpoints that help maintain occupant thermal comfort (23.5C-24C). As observed, a zone's temperature decreases ∼20 mins after a reduction in PDC setpoint, and vice versa. The table (illustrated in Table I) consists of different combinations of PDC setpoints (SP), weather data, ceiling temperature, and ambient table-level temperature data recorded by BLE-equipped temperature sensors (S_15, S_16, S_17). As described in Section IV-A, the weather values are quantized into 6 bins; for each combination of PDC setpoints, our lookup table consists of 20 mins data. While a 3-zone area conceptually contains $3^6 = 81$ unique tuples (each setpoint can have 6 values between (22.5°C,25°C), our training data was collected typically using 10 different combinations, which proved to be sufficient for effective PDC control.

### C. Real-Time Setpoint Control

We utilize the lookup table entries, together with input values for current weather and predicted occupancy (which together serve as a table index), to determine the "optimal" setpoint control values. For occupancy estimation, a state-of-the-art YoLov3 object detector [25] is executed on the camera images to obtain bounding boxes for human objects, which are then translated into real-world physical coordinates. For each matching (weather, occupancy) entry, we compute the CCF value by computing the cumulative deviation from the ideal ambient temperature (23.5°C for occupied; 25°C for unoccupied). Algorithm 1 details the CCF computation. Intuitively, the setpoint choice seeks to minimize the weighted sum of any deviations from the ideal ambient temperature. To emphasize user comfort, we apply a higher penalty (weight=2) if the projected temperature is higher than the desired value (i.e., if users will feel warmer); to concurrently prevent unnecessary cooling, we apply a lower penalty (weight=0.5) if the projected temperature is lower than the desired value.

**Algorithm 1** Trace-Based Lookup Algorithm

Input: Weather temperature (w), and Estimated Occupancy (o) for next 15-min epoch.
**for** j ∈ All Combination of PDC Setpoints (Weather=w,Occup=o) **do**
  **for** Minutes 1..20 **do**
    calculate CCF(j)=

$$\sum_{k \in zones} (2*(23.5-z_k^j)*\mathrm{I}(z_k^j > 23.5)\&\mathrm{I}(o_k^j \neq Unoccupied)+$$
$$0.5*(z_k^j - 23.5) * \mathrm{I}(z_k^j < 23.5)\& * \mathrm{I}(o_k^j \neq Unoccupied) +$$
$$0.5 * |z_k^j - 25.0| * \mathrm{I}(o_k^j = Unoccupied)$$

    where $z_k^j$ represents the predicted ambient temperature of the $k^{th}$ zone under the $j^{th}$ setpoint setting & I(...) represents the Indicator function.
  **end for**
**end for**
Find the combination which has the lowest cost–i.e., $\arg\min_j CCF(j)$

## VI. Experimental Results

We now quantify the performance of the proposed *OcAPO* approach. We test our approach in an uncontrolled (i.e., occupancy driven by natural usage) testbed setting. We compare our system to the baseline rule-driven, time-of-day based "BMS" (Building Management System) system. Both approaches were evaluated over 3 consecutive days (Wed, Thu, Fri), on separate weeks. Our evaluation focuses on (a) *OcAPO's* ability to maintain optimal thermal comfort during the deployment, (b) Occupancy Estimation errors, and (c) dynamic PDC setpoint control of *OcAPO* compared to the currently operating baseline.

### A. Thermal Comfort

We placed BLE sensors on tripod stands at the occupant level to measure thermal comfort. Fig 8 and 9 illustrates the zone and ceiling temperature recorded during the BMS PDC setpoint control and *OcAPO's* deployment phase, respectively, for Area-3 (Zone-11 and 12). Note that Zone_12 was never occupied in the deployment phase of both BMS and *OcAPO*. We calculate the comfort level only for Zone_11. We observe that *OcAPO* successfully maintains a thermal comfort of 23.3°C during the entire occupancy period. However, the BMS system does not maintain thermal comfort and is above the comfort level on average by 1.2°C. In Area-2, Zone 8 was never occupied. *OcAPO's* achieved a comfortable ambient temperature of 23.8°C in Zones 7 and 9 during the occupancy period, while BMS deviated from the ideal settings by an average of 1.1°C. In Area-1, however, both BMS and *OcAPO* were able to successfully assure thermal comfort with a deviation of 0.5°C and 0.3°C, respectively. This can be explained by noting that in Area 1 (unlike Areas 2 & 3), the readings of the ceiling-mounted thermal sensor (placed directly above the occupants) are consistent with the occupant-level ambient temperature (see Figure 6). Overall, *OcAPO* was able to maintain an average ambient temperature of 23.8°C under low, 23.8°C under medium, and 23.5°C for high occupancy conditions.

### B. Dynamic PDC setpoint Control

We also studied the behavior of dynamic PDC setpoint control (setting varying setpoints across zones for a given

TABLE I: Look-Up Table

| Time | Weather | SP_1 | SP_2 | SP_3 | O_1 | O_2 | O_3 | S_15 | S_16 | S_17 | CeilingTemp |
|---|---|---|---|---|---|---|---|---|---|---|---|
| T1 | 28 | 24.5 | 24 | 24 | LOW | HIGH | HIGH | 24.8 | 24.2 | 24.8 | 24.7 |
| T10 | 28 | 24.5 | 24 | 24 | LOW | HIGH | HIGH | 24.5 | 23.8 | 24.5 | 24.3 |
| T20 | 28 | 24.5 | 24 | 24 | LOW | HIGH | HIGH | 24.1 | 23.7 | 24.2 | 24.0 |

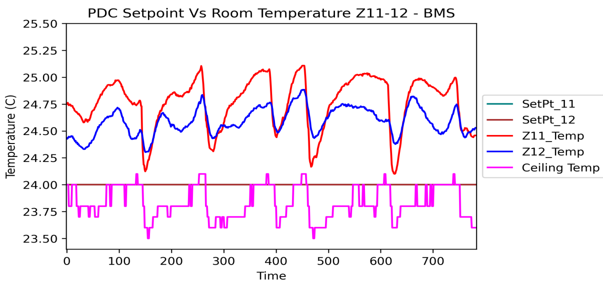


Fig. 8: Room Temperature Vs PDC Setpoint - BMS

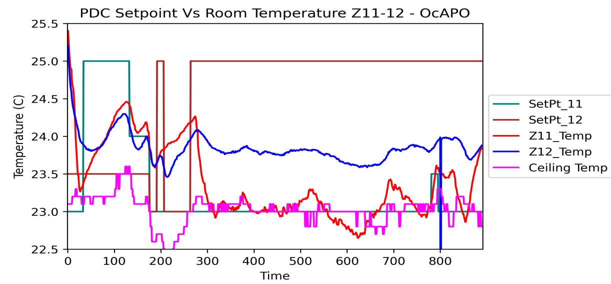


Fig. 9: Room Temperature Vs PDC Setpoint - OcAPO

area) compared to the occupancy-unaware static BMS control. For Areas 2, and 3, we observe that *OcAPO* set the PDC setpoint 1°C *above* the BMS setpoint of 24°C when the zone is not occupied and 1-1.5°C *below* the BMS setpoint when the zone is occupied. *OcAPO* effectively executes dynamic setpoint control to assure occupant thermal comfort. In the unoccupied zones, *OcAPO* set the PDC setpoint at 25°C–see, for example, Zone-12 in Figure 9.

### C. Occupancy Estimation

State-of-the-art object detectors experience errors (a) in partially-occluded environments, (b) for more distant objects that are smaller in size; such errors can cause *OcAPO* to perform incorrect PDC setpoint control. While the the camera-based occupancy estimator had an accuracy of 93.7% (in classifying occupancy across the 4 bins), the use of temporal smoothing (over multiple readings within a 15-minute window) helped reduce the estimation error by 23.8%. Overall, over a 6-hour operational period, occupancy estimation errors caused *OcAPO* to specify incorrect PDC setpoints thrice (i.e., for a total of 45 mins).

## VII. Discussion and Future Work

**Accommodating Transient Artifacts:** *OcAPO* relies on the data collected over a designated period to adjust the PDC setpoints. However, the relationship between PDC setpoint and table-level ambient temperature can experience short-term deviations, due to usage artifacts (e.g., users operating a heat-generating GPU server), which *OcAPO* currently is incapable of tackling. Adapting to such localized "heat island" effects requires additional and extensive sensor instrumentation, and is also constrained by the longer response transient of PDC.

**Cost and Scalability:** *OcAPO* requires the creation of the look-up table, using data collected to capture seasonal variations. While such variation is minimal in Singapore's tropical climate, the observation period will need to be much longer (even spanning a year) for environments with greater seasonal fluctuations. *OcAPO's* computational overhead arises primarily from executing the DNN models for real-time occupancy estimation, and can be reduced by lowering the estimation frequency (e.g., to once every 5 minutes). Scaling *OcAPO* to all open-floor spaces on our campus will require more intensive data collection, with 6 person-month effort, across 120 camera feeds. Supporting *OcAPO* will, however, require first require capital investment to install PLC-controllable PDC systems across all such spaces.

**Modelling Historical Time-Varying Changes:** *OcAPO* currently performs adaptation based on a model of ambient temperature evolution over a 20 min window, and is thus oblivious to the longer term ambient temperature readings and weather conditions. In reality, it is likely that adaptation polices that incorporate longer-term memory (e.g., 12 hour history) might offer more optimal energy-vs.-comfort tradeoffs. However, developing such "trajectory-based models" is hard and requires significantly larger volumes of longitudinal training data.

## VIII. Conclusion

Motivated by empirical observations of a material, occupancy-dependent discrepancy between the temperature recorded by ceiling-mounted sensors (which drives the PDC control loop) and occupant-level ambient temperature in an open plan collaborative $\sim$1000m$^2$ workspace, we have introduced *OcAPO*, an occupancy-aware, adaptive PDC setpoint control system. *OcAPO* takes real-time and diurnal occupancy estimates (obtained automatically by analyzing camera feeds), together with current weather conditions, as input and and utilizes a trace-based lookup table to dynamically specify an optimum tuple of PDC setpoint values for a given area. A real-world deployment of *OcAPO* establishes that it can maintain an optimum thermal comfort (range: $23.5 - 24.0$°C ) with a tolerance of 0.2°C, compared to a conventional BMS system (where the setpoint is always maintained at 24.0°C) which exceeds the optimal thermal comfort by $\geq$ 0.8°C. Overall, *OcAPO* is able maintain an average ambient temperature of 23.8°C under low, 23.8°C under medium, and 23.5°C under high occupancy conditions.

## Acknowledgment

This work was supported partially by the National Research Foundation, Singapore under its NRF Investigatorship grant

(NRF-NRFI05-2019-0007) and by the BCA GBIC grant. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

## References


[1] M. González-Torres, L. Pérez-Lombard, J. F. Coronel, I. R. Maestre, and D. Yan, "A review on buildings energy information: Trends, end-uses, fuels and drivers," *Energy Reports*, vol. 8, pp. 626–637, 2022.

[2] F. Oldewurtel, D. Sturzenegger, and M. Morari, "Importance of occupancy information for building climate control," *Applied Energy*, vol. 101, pp. 521–532, 2013.

[3] S. Azizi, R. Rabiee, G. Nair, and T. Olofsson, "Effects of positioning of multi-sensor devices on occupancy and indoor environmental monitoring in single-occupant offices," *Energies*, vol. 14, no. 19, 2021.

[4] M. Gruber, A. Truschel, and J.-O. Dalenback, "Co2 sensors for occupancy estimations: Potential in building automation applications," *Energy and Buildings*, vol. 84, pp. 548–556, 2014.

[5] N. Li, Z. Yang, B. Becerik-Gerber, and M. Orosz, "Towards Energy Savings From a Bimodal Occupancy Driven Hvac Controller in Practice," *Proceedings of the CIB W78 2013: 30th International Conference –Beijing, China, 9-12 October*, no. Doe 2011, pp. 680–689, 2013.

[6] S. R. West, J. K. Ward, and J. Wall, "Trial results from a model predictive control and optimisation system for commercial building HVAC," *Energy and Buildings*, vol. 72, pp. 271–279, 2014.

[7] T. C. Seng and T. Soe, "Getting nature to help in energy efficiency of the air conditioning system," *Energy Procedia*, vol. 143, pp. 230–236, 2017.

[8] C. Nesler, "Adaptive control of thermal processes in buildings," *IEEE Control Systems Magazine*, vol. 6, no. 4, pp. 9–13, 1986.

[9] B. Dong, K. P. Lam, and C. P. Neuman, "Integrated building control based on occupant behavior pattern detection and local weather forecasting," *Proceedings of Building Simulation 2011: 12th Conference of International Building Performance Simulation Association*, vol. 3, pp. 193–200, 2011.

[10] T. Kitzberger, J. Kotik, and T. Pröll, "Energy savings potential of occupancy-based HVAC control in laboratory buildings," *Energy and Buildings*, vol. 263, p. 112031, 2022.

[11] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, and Y. Agarwal, "Sentinel: Occupancy based HVAC actuation using existing wifi infrastructure within commercial buildings," *SenSys 2013 - Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, 2013.

[12] S. Nagarathinam, H. Doddi, A. Vasan, V. Sarangan, P. Venkata Ramakrishna, and A. Sivasubramaniam, "Energy efficient thermal comfort in open-plan office buildings," *Energy and Buildings*, vol. 139, pp. 476–486, 2017.

[13] J. R. Dobbs and B. M. Hencey, "Predictive HVAC Control Using a Markov Occupancy Model," in *Proceedings of the American Control Conference*, no. June, 2014, pp. 1057–1062.

[14] J. Shi, N. Yu, and W. Yao, "Energy Efficient Building HVAC Control Algorithm with Real-time Occupancy Prediction," *Energy Procedia*, vol. 111, no. September 2016, pp. 267–276, 2017.

[15] J. Hou, H. Li, N. Nord, and G. Huang, "Model predictive control under weather forecast uncertainty for HVAC systems in university buildings," *Energy and Buildings*, vol. 257, p. 111793, 2022.

[16] "Data predictive control using regression trees and ensemble learning," *2017 IEEE 56th Annual Conference on Decision and Control, CDC 2017*, vol. 2018-January, no. September, pp. 4446–4451, 2017.

[17] F. Smarra, A. Jain, T. de Rubeis, D. Ambrosini, A. D'Innocenzo, and R. Mangharam, "Data-driven model predictive control using random forests for building energy optimization and climate control," *Applied Energy*, vol. 226, no. February, pp. 1252–1272, 2018.

[18] T. Wei, Y. Wang, and Q. Zhu, "Deep Reinforcement Learning for Building HVAC Control," *Proceedings - Design Automation Conference*, vol. Part 12828, no. 2, 2017.

[19] G. Gao, J. Li, and Y. Wen, "DeepComfort: Energy-Efficient Thermal Comfort Control in Buildings Via Reinforcement Learning," *IEEE Internet of Things Journal*, vol. 7, no. 9, pp. 8472–8484, 2020.

[20] M. Biemann, F. Scheller, X. Liu, and L. Huang, "Experimental evaluation of model-free reinforcement learning algorithms for continuous HVAC control," *Applied Energy*, vol. 298, no. March, p. 117164, 2021.

[21] "MARCO - Multi-Agent Reinforcement learning based COntrol of building HVAC systems," *e-Energy 2020 - Proceedings of the 11th ACM International Conference on Future Energy Systems*, no. Ml, pp. 57–67, 2020.

[22] M. Esrafilian-najafabadi and F. Haghighat, "Energy Buildings Occupancy-based HVAC control using deep learning algorithms for estimating online preconditioning time in residential buildings," *Energy Buildings*, vol. 252, p. 111377, 2021.

[23] D. Azuatalam, W. L. Lee, F. de Nijs, and A. Liebman, "Reinforcement learning for whole-building HVAC control and demand response," *Energy and AI*, vol. 2, p. 100020, 2020.

[24] B. Chen, Z. Cai, and M. Bergés, "Gnu-RL: A precocial reinforcement learning solution for building HVAC control using a differentiable MPC policy," *BuildSys 2019 - Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pp. 316–325, 2019.

[25] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *CoRR*, vol. abs/1804.02767, 2018.