Inverted HVAC: Greenifying Older Buildings, One Room at a Time

SAMAR ABBAS, ABU BAKAR, YASRA CHANDIO, and KHADIJA HAFEEZ, Department of Computer Science, LUMS, Pakistan
AYESHA ALI, Department of Economics, LUMS, Pakistan
TARIQ M. JADOON, Department of Electrical Engineering, LUMS, Pakistan
MUHAMMAD HAMAD ALIZAI, Department of Computer Science, LUMS, Pakistan

26

Emerging countries predominantly rely on room-level air conditioning units (window ACs, space heaters, ceiling fans) for thermal comfort. These distributed units have manual, decentralized control leading to sub-optimal energy usage for two reasons: excessive setpoints by individuals and inability to interleave different conditioning units for energy savings. We propose a novel *inverted* HVAC approach: cheaply retrofitting these distributed units with "on-off" control and providing centralized control augmented with room and environmental sensors. Our binary control approach exploits an understanding of device consumption characteristics and factors this into the control algorithms to reduce consumption. We implement this approach as HAWADAAR in a prototype 180ft^2 room to evaluate its efficacy over a 7-month period experiencing both hot and cold climates. Through a post analysis, we show that our on-off algorithms are not far from a theoretically optimal approach based on *a priori* information that precisely knows the optimal control points to minimize consumption. We collect enough evidence to plausibly scale our empirical evaluation, demonstrating countrywide benefits: with just 20% market penetration, HAWADAAR can save up to 6% of electricity per capita in residential and commercial sectors—resulting in a substantial countrywide impact.

CCS Concepts: • **Computer systems organization** → *Sensors and actuators*;

Additional Key Words and Phrases: Thermal comfort, energy efficiency, inverted HVAC

ACM Reference format:
Samar Abbas, Abu Bakar, Yasra Chandio, Khadija Hafeez, Ayesha Ali, Tariq M. Jadoon, and Muhammad Hamad Alizai. 2018. Inverted HVAC: Greenifying Older Buildings, One Room at a Time. *ACM Trans. Sen. Netw.* 14, 3–4, Article 26 (November 2018), 26 pages.
https://doi.org/10.1145/3229063

1 INTRODUCTION

We are in an era of global warming, with most studies indicating an excessive use of energy being its leading cause. By most accounts, the greatest contribution to global energy expenditure and

Authors' addresses: S. Abbas, A. Bakar, Y. Chandio, and K. Hafeez, Department of Computer Science, LUMS, DHA, Lahore, Punjab, 54792, Pakistan; emails: {16030057, abubakar, yasra.chandio, 15030019}@lums.edu.pk; A. Ali, Department of Economics, LUMS, DHA, Lahore, Punjab, 54792, Pakistan; email: ayeshaali@lums.edu.pk; T. M. Jadoon, Department of Electrical Engineering, LUMS, DHA, Lahore, Punjab, 54792, Pakistan; email: jadoon@lums.edu.pk; M. H. Alizai (corresponding author), Department of Computer Science, LUMS, DHA, Lahore, Punjab, 54792, Pakistan; email: hamad.alizai@lums.edu.pk.

1550-4859/2018/11-ART26 $15.00
https://doi.org/10.1145/3229063

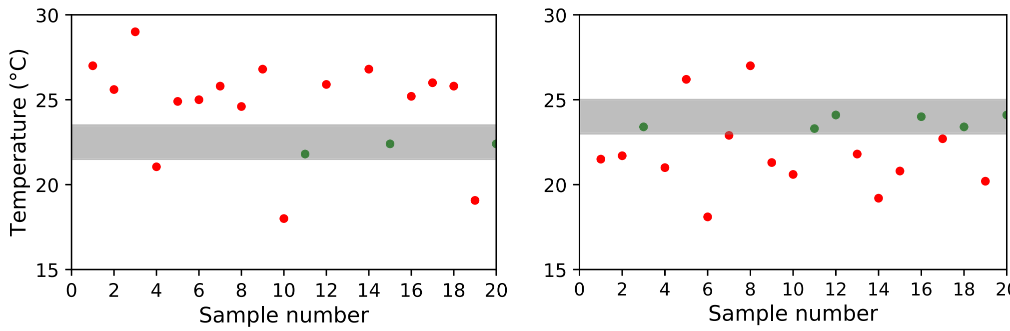

Fig. 1. Temperature measurements from randomly sampled rooms in the administrative and residential complexes of our university with room-level conditioning units: Temperatures often kept too high in winters or too low in summers with neither the incentive nor the capability to achieve thermal comfort at low energy budget. Our goal is to push these red dots into the shaded region (recommended range). Energy saving is a reward.

the ensuing green house gasses will come from emerging Asian countries [13, 26], such as China, India, Pakistan, and Bangladesh that in combination account for nearly half of the world's population [14]. With their economic upsurge, the energy demands of these countries are rapidly rising; thermal comfort of built spaces making a significant proportion. It is projected that there will be a tenfold increase in the world consumption of energy for cooling by 2050 [23]. Thus, for example, China alone is expected to surpass the USA by 2020 as the world's biggest consumer of electricity for air conditioning estimated at a trillion kilowatt hours (kWh) [12].

With the late uptake of efficient thermal comfort systems (like HVAC) in these regions, the majority of buildings still employ room-level units—such as window or split ACs, space heaters, ceiling, and sometimes ventilation fans (see Figure 3 for a typical room)—for thermal comfort. These buildings will be maintained for several decades with at least 80% to last beyond 2050 [22, 28], considering housing needs, economic constraints, as well as heritage protection. The challenge is compounded with the continental climate in most countries in the region, which is characterized by extreme temperature variations, both daily and seasonally. Consequently, this extreme climate also results in *excessive* energy usage.

Anecdotally, room occupants set these distributively controllable units to exceed recommended temperature settings creating the kind of indoor environment in which occupants wear sweaters and use blankets in July [11]. We validate this observation by a survey of temperature readings shown in Figure 1; the data validating that a distributive approach to temperature setting results in inefficiency. These aggressive setpoints stem from (i) a psychological reaction to outside temperatures (which can be extreme), (ii) possibly inappropriate AC sizing for the room resulting in insufficient comfort at a person's location in the room, or (iii) a lack of per device thermal control (space heater without thermostat). Furthermore, these excessive settings also fail to meet conditioning standards [4], thus being detrimental for the health of the individual.

We propose a novel approach to solve this geographically unique problem; *distributed room-level conditioning* governed by a standards-compliant control abstraction, the *centralized building-level thermostat*. This is achieved by adding "smartness" to existing devices to maintain thermal comfort whilst saving energy. This approach is interesting as an inversion to the HVAC approach to managing thermal comfort: a *centralized conditioning* unit, the AHU, that distributes air to zones controlled by individual thermostats.

We extend HAWADAAR[1] [20], as an implementation of this inverted HVAC approach to evaluate its efficacy. HAWADAAR has three novel aspects that set it apart from the existing literature. First, its ability to interleave several modes of achieving thermal comfort—such as cooling using an AC, an evaporative cooler, or through air circulation using a fan—with an objective to minimize electricity consumption. Second, its *adaptive* two-position control strategy that intelligently orchestrates these units, accounting for thermal impacting factors such as internal and external thermal loads, room insulation, and device characteristics (transients and short cycling). Third, its ability to handle a wide variety of heating and cooling devices to deliver thermal comfort across a wide range of weather conditions 24/7 all year around.

Our work, thus, has four significant contributions.[2]

Inverted HVAC: We propose an inverted HVAC approach and its IoT-inspired architecture in Section 2. This approach to implementing thermal comfort is novel and especially pertinent to the socioeconomic and climatic constraints of emerging countries. We introduce intelligent control algorithms for the centralized thermostat. These algorithms are based on an empirical and theoretical understanding of the (conditioning) device constraints (Section 3), as well as factors impacting thermal conditioning, such as modes of heat transfer and perceived thermal comfort (Section 4).

System Evaluation: We empirically evaluate a pilot implementation of this approach—HAWADAAR—over a 7-month period to demonstrate (a) the efficiency of our algorithms in achieving setpoints with ±0.5°C and ±1°C tolerance, and (b) energy savings (at <50% duty cycle), in Section 5.

Post Analysis: We derive a thermal model of air conditioning devices based on a post analysis of the empirical data. Since this model is based on complete information about the heating/cooling characteristics of devices under varying weather conditions, it precisely knows the optimal control points for on/off to minimize the duty cycle. We compare the simulated results based on this thermal model with HAWADAAR to show how far is our practical implementation from a theoretically-optimal control approach.

Country-scale Forecasting: Finally, we extend our experimental results to predict country-scale energy savings in Section 8. We arrive at a conservative estimates of (6% per capita) projected savings for a given market penetration (20%) of a HAWADAAR-like system.

We discuss related work along with future outlook in Section 8 before concluding the article in Section 9.

2 HAWADAAR: THE INVERTED HVAC

We now describe HAWADAAR in detail to elaborate the inverted HVAC approach and its IoT-inspired architecture. We will then present setup details of a prototype system for a single room.

2.1 What's an Inverted HVAC?

The *inversion* of HVAC approach in our proposed system stems from the inversion of the location of control and conditioning units. We propose to use a set of disparate and distributed conditioning units to enforce control from a central location for an alternate to modern HVAC systems. This approach makes sense only when viewed in the context of the socioeconomic background we advocated earlier: populous emerging economies with widespread installation of room-level units

[1]Local slang for *ventilated and (air) conditioned space*.

[2]Parts of this work were presented in Reference [20]. Our post analysis in Section 6 is entirely new and provides deeper insights.

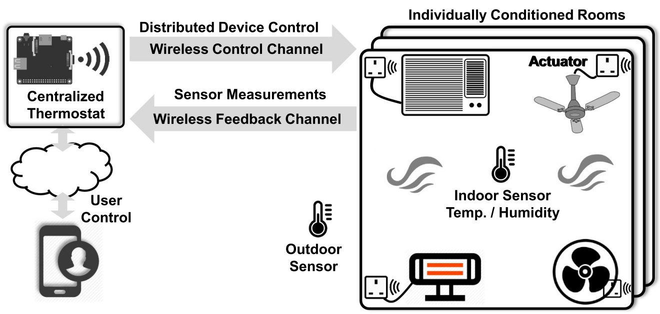

Fig. 2. Inverted HVAC architecture: IoT retrofitting for distributed conditioning via a centralized control abstraction, to improve thermal comfort and energy efficiency of legacy buildings.

where fitting HVAC is cost prohibitive. The novelty of our approach lies in identifying this unique opportunity.

The implementation of this approach, as shown in Figure 2, extends an IoT-inspired architecture that involves augmenting every installed unit[3] with wireless on-off control and distributively sensing temperature and humidity through a wireless back-channel (direct or multi-hop). A centralized hub-like device hosts a software-based control abstraction, the *centralized thermostat* (CT), to trigger room-local control for the entire building. The CT can be configured on a per room basis by a user (e.g., the building owner), through a smartphone app, to either deliver a *setpoint* or *personalized comfort*—based on ASHRAE's personal comfort metric (PMV) [4]—within respective, standards-compliant comfort bounds. Further indoor sensors can be added to improve the per room sensing reliability and coverage; however, we emphasize on minimizing retrofitting costs assuming that the indoor sensor is deployed at a pertinent location where the required thermal objective has to be achieved. Additionally, an outdoor sensor is required for estimating a room's heat transfer coefficient (see Section 4.3) for intelligent actuation. For global access and refactoring into other applications (e.g., smart grids), the CT can be hosted in the cloud where an appropriate API layer can expose its control and measurements to user applications.

2.2 A Prototype Implementation of Hawadaar

We now present a practical realization of the inverted HVAC—Hawadaar—to demonstrate and then evaluate the efficacy of this approach. Figure 3 shows our deployment in a 180ft^2 room equipped with multiple air conditioning units (fan, AC, space heater). This room is representative of closed spaces in the developing world, where all or a subset of these appliances are present. Our room-level evaluation can thus be extrapolated to homes, apartments, hotels, office buildings, and so on. Furthermore, this room has three external facing walls and roof exposed to the elements, representing a challenging scenario for conditioning the room.

We reduce the retrofitting complexity by choosing to control each device with a COTS[4] smart-plug [35], which costs as low as \$2 each when ordered in bulk. The existing device sockets are inserted into these plugs.[5] For convenience, we choose Z-wave-based plugs [35] and temperature and humidity sensors [36]. We emulate our centralized thermostat using a Z-wave dongle attached to a RPi with all our control algorithms implemented on it. Thus, depending upon the magnitude

[3]ACs, heaters, ceiling and ventilation fans, personal comfort devices.
[4]Commercial off-the-shelf.
[5]AC might additionally require a relay in between to cater for high surge currents.

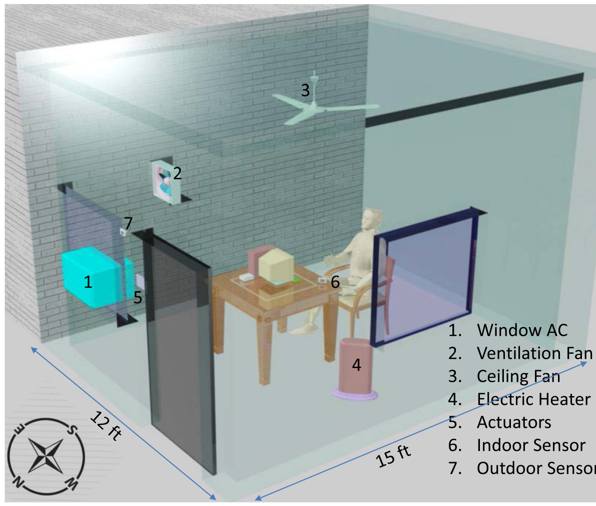

Fig. 3. HAWADAAR Deployment setup: A typical office room (180ft^2) with legacy devices (1–4) and IoT retrofitting (5–7). The roof and three walls are exposed to elements.

of supply order, the current retrofitting cost per room with three conditioning units could be as low as $10, including smart-plugs, sensor, and the amortized cost of CT.

We choose an arbitrary location for the sensor, i.e., the wall opposite to the AC when doing setpoint-based control, similarly on the desk next to the fictional occupant in Figure 3 for PMV-based control. This sensor, being wireless, can be placed at any appropriate location and, thus, the above evaluation is sufficient. These (one-per-room) sensors report their values once per minute to the CT.

We note that our plug-based retrofitting limits our control to simple on-off for each device and can introduce issues with regards to device safety if the control algorithm does not have appropriate hysteresis. Entirely for this purpose, we evaluate device constraints in the next section and use them to tune our algorithms in Section 4 to prevent energy waste as well as wear-and-tear.

3 DEVICES AND CONSTRAINTS

We first highlight thermal characteristics of air conditioning devices under consideration, then briefly review what device constraints are relevant and how they impact control, and subsequently derive these constraints empirically.

3.1 Air Conditioning Units Under Consideration

Figure 4 depicts the three devices currently employed by HAWADAAR for distributed conditioning. With a 1-ton cooling capacity, the single-unit window AC provides *convective* cooling through a single-stage heatpump. In our setup, we position the setpoint of the AC's onboard thermostat to its minimum value, thus allowing HAWADAAR to independently control the AC (without altering the control circuitory) for setpoints above this minimum value. The electric heater implements *radiant* heating through its three halogen elements, each consuming 400W. When tested under cooling and heating loads greater than 10°C, the minimum and maximum room temperatures achieved through the AC and heater are 18°C and 26°C, respectively. Both these values fall well outside the comfort zones for the respective seasons (see Section 5.2.2), thereby, already highlighting the potential of energy conservation through duty cycling; more so in the case of an electric heater

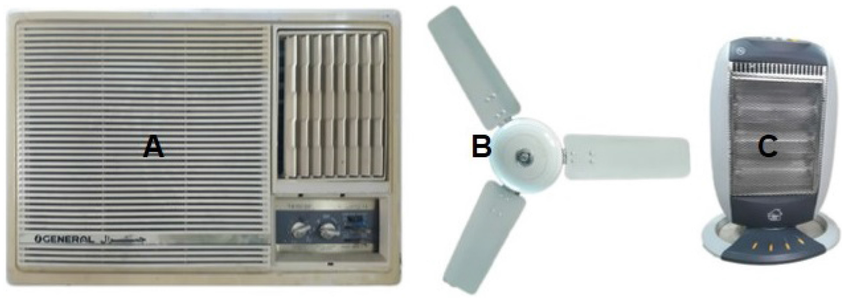

Fig. 4. Air conditioning units under consideration: (A) single unit window AC with 1-ton cooling capacity; (B) ceiling fan with sweep size of 36in.; (C) electric space heater with three halogen elements.

without thermostat. The ceiling fan has a sweep size of 36in. and can be regulated at five different speeds, with a maximum air velocity of 3.77ms^{-1}. We set the fan speed to its maximum for on-off control by Hawadaar. While the fan only provides air movement, the AC and heater impact both temperature and humidity.

3.2 Relevant Constraints and Their Impact

How accurately can Hawadaar maintain a setpoint? This depends on how aggressively it can actuate the respective air conditioning devices. For example, in cooling, two-position control is achieved by switching a device *on* at temperatures exceeding T_{on} and *off* at temperatures below T_{off}, and ideally T_{on} should be identical to T_{off}. However, such a strict thermal objective is not achievable due to *deadband* requirement to prevent repeated on-off cycles, as well as two device-specific constraints: *transient power* (the increased power consumption at startup) and *short cycling* (actuating it faster than the specified rate). Hawadaar must be cognizant of these *hard* constraints; otherwise, the former could increase the overall energy consumption of the system and the latter could potentially damage or reduce the life span of a device. Thus, to calibrate Hawadaar's control algorithm, we study both these characteristics in our setup and calculate the smallest safe interval t_{safe}, the minimum duration between switching the device off and then on. Alternatively, one could pick a large enough t_{safe} to make the algorithm agnostic to such device constraints, but this could compromise thermal comfort [5]. The idea is to set t_{safe} to the larger of two intervals, i.e., $max(t_{tp}, t_{sc})$, where t_{tp} and t_{sc} refer to the respective transient power and short cycling constraints.

3.3 Deriving Hard Device Constraints

Figure 5 shows the power consumption of devices in transient and steady state. To avoid the penalty of transient power due to frequent switching, the minimum duration (i.e., t_{tp}) for which a device must remain off is defined by the interval over which the excess transient power equals the power consumed if the device had not been switched off. Thus, we need to switch off for at least long enough that this saved steady-state power compensates for the high transient power at the next startup. A longer off duration will indeed translate into power gains. In Figure 5, based on the area under the curve, we compute t_{tp} simply as $t_{tp} = \frac{E}{P}$, where E is the excess transient energy and P is the steady-state power consumption. With regard to short cycling, there are no theoretical restrictions on the switching frequency of electric heaters and fans. Whereas, for window ACs, manufacturers specify a minimum off-duration of 3min [1, 17]. Table 1 enumerates t_{safe} for each device that serves as a *hard* input constraint on the maximum switching frequency of the

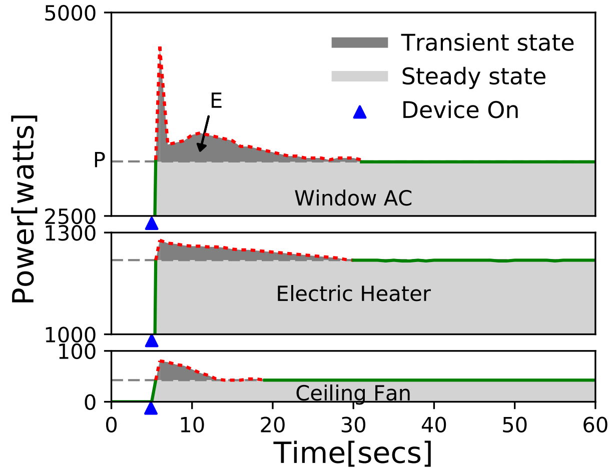

Fig. 5. Transient vs. steady-state power consumption.

Table 1. Device Constraints: t_{tp} is Transient Power, t_{sc} Short Cycling Constraint, and t_{safe} is Larger of These Two

Device	**t_{tp} (s)**	**t_{sc} (s)**	**$t_{safe} = \max(t_{tp}, t_{sc})$**
AC	1.62	180	180
Heater	0.63	—	0.63
Ceiling Fan	4.3	—	4.3

control algorithm. However, as we notice in Section 5.2.1, this switching frequency is not actuated even for very narrow comfort bounds.

4 THE CENTRALIZED THERMOSTAT

With device constraints clearly defined, we are now ready to build a control abstraction, i.e., the centralized thermostat. The CT allows to configure per room thermal settings through a mobile-app, to either maintain a setpoint or deliver personalized comfort (PMV), within standards-compliant tolerance. First, we explain why a reactive control strategy is employed in HAWADAAR instead of model predictive control (MPC) [9]. We then describe our control algorithm and its heat-transfer prediction mechanism, which is needed for the dynamic adaption of comfort bounds (i.e., the tolerance range) based on thermal load.

4.1 Why Two-Position Control?

The CT runs an intelligent control algorithm that performs *adaptive* two-position control: dynamically adjusts the *soft* constraint (comfort bounds) around a fixed setpoint to satisfy the *hard* device constraints. We use this simple reactive control strategy instead of a complex model predictive approach for three key reasons: First, unlike centralized HVAC taking tens of minutes to take effect, these room-level conditioning units affect human comfort immediately [18]. Second, room insulation levels in older buildings are suboptimal, thereby requiring an immediate response to deliver adequate thermal comfort. Finally, a reactive control strategy is inherently sensitive to changing thermal loads; for example, air conditioning will run for longer if there are more occupants than

ALGORITHM 1: Setpoint-based adaptive two-position control.

Input: desired setpoint (T_s), tolerance

while True **do**
    $T_{in} \leftarrow$ read_sensor()
    $T_{on} \leftarrow T_s + \frac{\text{tolerance}}{2}$
    $T_{off} \leftarrow T_s - \frac{\text{tolerance}}{2}$
    if $T_{in} > T_{on}$ **then**
        switch_on(AC)
    end
    else if $T_{in} \leq T_{off}$ **then**
        predicted_temperature $\leftarrow T_{t_{off}+t_{safe}}$
        if predicted_temperature $\leq T_{on}$ **then**
            switch_off(AC)
            update(prediction_parameters)
            reset(tolerance)
        end
        else
            extend_tolerance(tolerance)
        end
    end
end

usual in the room. This eliminates the need for complex thermal load estimations [29], such as human occupancy, further simplifying our pilot system design.

4.2 Algorithm: Adaptive Two-Position Control

We first describe the algorithm in the context of maintaining a setpoint and later extend our description to delivering personalized comfort. We do not make any assumptions regarding the relative humidity and use *heat index* (a.k.a. "apparent temperature" or "feels like") as our temperature metric. Thus, a setpoint is defined in terms of heat index, not the ambient temperature. Hence, before making a control decision, the temperature value is first converted into its corresponding heat index (HI), which is calculated as follows:

$$\begin{aligned} HI(T,R) = c_1 + c_2T + c_3R + c_4TR + c_5T^2 \\ + c_6R^2 + c_7T^2R + c_8TR^2 + c_9T^2R^2, \end{aligned} \tag{1}$$

where T is the air temperature, R is relative humidity, and c_n are Rothfusz regression constants [10]. For brevity, we use the general term "temperature" (instead of heat index) to simplify the description of our algorithm.

4.2.1 Setpoint Algorithm. To begin, the algorithm needs two inputs, the required setpoint T_s and comfort bounds (soft constraint); the latter defines two control positions T_{off} and T_{on} symmetrically around T_s. Thus, with a cooling device, to achieve T_s as the average temperature, we must satisfy $T_{t_{off}+t_{safe}} \leq T_{on}$, i.e., once the device is turned off, the temperature after the minimum safe off duration t_{safe} (hard constraint) will not exceed T_{on}. This may, depending upon the thermal load, require the algorithm to symmetrically extend the tolerance range on both sides of T_s, as described in Algorithm 1. For this, the algorithm needs to predict at time t_{off} (cf. Secton 4.3) the temperature at time $t_{off} + t_{safe}$ to decide whether to turn off the device or extend the tolerance range to meet

ALGORITHM 2: PMV-based two-position control.

```
Input:  comfort_range, set of on devices (ON)
while True do
    calculate(PMV)                                  // see Equation (2)
    if PMV > upper_bound then
        if ON = ∅ then
            v ← AIR_VELOCITY_fan
            calculate(PMV)
            if PMV ∈ comfort_range then
                ON ← ON ⊎ fan                       // switch-on fan
            end
            else
                v ← AIR_VELOCITY_AC
                ON ← ON ⊎ AC                        // switch-on AC
            end
        end
        else if fan ∈ ON then
            ON ← ON \ fan                           // switch-off fan
            v ← AIR_VELOCITY_AC
            ON ← ON ⊎ AC
        end
    end
    else if PMV ≤ lower_bound then
        ON ← ∅                                      // switch-off all devices
    end
end
```

the hard constraint. With a minor adjustment, i.e., by swapping T_{off} with T_{on}, this algorithm is also used to actuate a heating device.

4.2.2 PMV Algorithm. The same baseline algorithm is used for delivering personalized comfort with the following two extensions: First, the thermal constraint is not the heat index but PMV described as follows:

$$PMV = f(M, T_a, T_r, v, P_a, I_{cl}), \tag{2}$$

where M is the metabolic rate of the occupant (assumed $70W/m^2$ for an office worker); T_a is the air temperature; T_r is the mean radiant temperature (set equal to T_a); v is the relative air velocity in m/s^{-1}; P_a is the relative humidity; and I_{cl} is the clothing insulation factor of the occupant (set to 0.6 *clo* assuming a usual office dress code of a long sleeved shirt with trousers). In our work, we assume some of these parameters to calculate PMV and realize that in practice these are difficult to accurately ascertain; however, this does not affect the fidelity of the results. Second, while the heat index-based setpoint only utilizes the heater and AC, additional factors in Equation (2), such as air movement, also allow us to use the ceiling fan for maintaining a desired PMV level. Thus, as described in Algorithm 2, we program the CT to prioritize the use of low-energy fan, and only turn on the AC when air circulation alone cannot keep the PMV within the required comfort bounds.

4.3 Predicting the Room's Heat Transfer Rate

The room's heat transfer depends upon both external and internal thermal loads. The external thermal loads result in heat transfer through the building envelope from external elements such

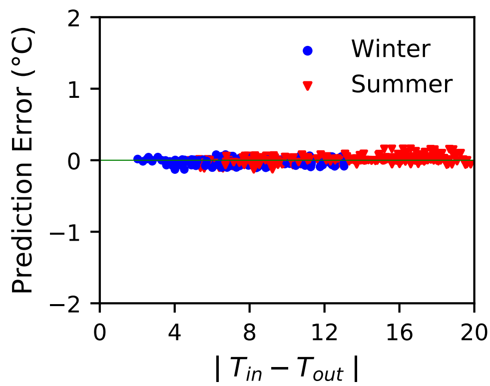

Fig. 6. Prediction accuracy: Our simple model achieves high accuracy with a root-mean-square error of 0.18°C.

as the sun, the earth, and the outside environment. While, internal thermal loads come from heat generated within the room by people, lighting, equipment, and so on. As discussed in the previous section, to satisfy the hard constraint, the CT needs to predict the temperature $T_{t_{off}+t_{safe}}$. The prediction model should thus account for both external and internal thermal loads whilst being simple and self calibrating. We note that the maximum duration of this prediction, in our setup, is just three minutes imposed by the AC. This, by the way, also provides a maximum theoretical duration of discomfort. Since we dynamically update our prediction parameters, as discussed below, our system can repeatedly fix prediction errors.

Instead of developing a complex thermal model for the room, we use realtime sensor measurements to glean the heat transfer coefficient using Newton's law of cooling as a first approximation. A similar strategy has also been employed in a personalized comfort system [18]. According to the law, "the rate of heat loss of a body is proportional to the difference in temperatures between the body and its surroundings." Thus, given an outside temperature T_{out} and room temperature T_{in}, the rate of thermal energy loss for the room is proportional to the temperature difference:

$$\frac{dT}{dt} = -k(T_{in} - T_{out}). \tag{3}$$

We set $T = T_{t_{off}+t_{safe}}$, $T_{in} = T_{off}$, and $t = t_{safe}$ and solve the above equation to estimate a room's heat transfer coefficient (k):

$$k = \frac{\ln \frac{(T_{t_{off}+t_{safe}} - T_{out})}{(T_{off} - T_{out})}}{t_{safe}}. \tag{4}$$

However, we repeatedly update k just before T_{on} (i.e., when all the devices are off), aiming for an *aggregate account of both external and internal thermal loads* using a single heat transfer coefficient (k). The sensed $T_{t_{off}+t_{safe}}$ in round n is used to calculate the k value to predict $T_{t_{off}+t_{safe}}$ in round $n+1$ as follows:

$$T_{t_{off}+t_{safe}} = T_{out} + (T_{off} - T_{out})e^{-kt_{safe}}. \tag{5}$$

Figure 6 depicts the accuracy of our prediction mechanism for numerous samples over a wide range of indoor and outdoor temperature differences. We can easily conclude that this model is sufficiently accurate for a reactive control strategy with a maximum required prediction length of just 3min.

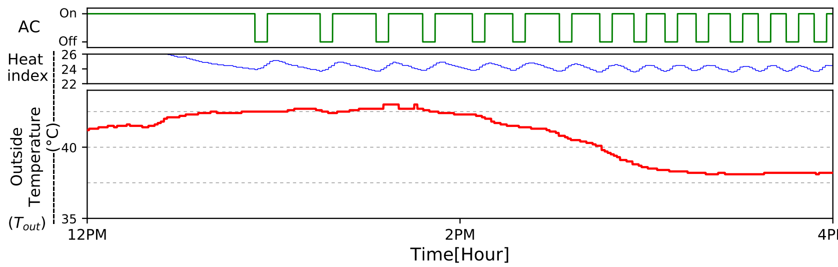

Fig. 7. The best effort service: HAWADAAR successfully adapts the thresholds of its two-positions symmetrically oscillating around the setpoint. The heat index reflects the room temperature.

5 EVALUATING THE DEPLOYMENT

Our pilot deployment seeks answers to two fundamental questions regarding a thermal comfort solution: (i) *How comfortable is it*? (ii) *What are its energy benefits*? To answer the first question, we evaluate the minimum temperature tolerance (*best effort* service) required to operate HAWADAAR. This best effort service is also relevant to satisfy high comfort requirements of a demanding user and to stress test Algorithm 1. To answer the second question, we compare energy consumption in multiple settings. For example, when varying tolerance around a fixed setpoint; when exceeding appropriate temperature settings, based on our anecdotal observation substantiated in Figure 1, and by interleaving devices of variable energy consumption, as in Algorithm 2.

Although our deployment setup is in place since November 2016, here we only report results from experiments during two challenging weather spells occurring between January 25 and 28, 2017, and May 1 and 15, 2017, respectively. Throughout these experiments, the room occupation varied between 0 (at night) to at most 3 (during the day) occupants.

5.1 Thermal Comfort: The Best Effort Service

This part of the evaluation corresponds to the adaptive control that minimizes the tolerance range ($|T_{off} - T_{on}|$) whilst satisfying hard device constraints. Thus, as described in Algorithm 1, the off-time of a device is fixed at t_{safe}, while the on-time (duty cycle) is a function of thermal load at a certain instant of time. We expect the algorithm to minimize its tolerance range, given the device constraint (t_{safe}), such that $\frac{T_{off}+T_{on}}{2}$ is always approximately equal to the required setpoint (T_s). In other words, we want our system to achieve two goals: (i) dynamically adapt the tolerance range and thus the duty cycle based on the thermal load, and (ii) deliver an average temperature—in terms of heat index—that does not deviate from the required setpoint. An inaccurate temperature prediction at T_{off} could potentially lead to such deviations. Figure 7 verifies the fidelity of our algorithm and its temperature prediction under varying but high external thermal loads, when the setpoint is kept near the middle (24°C) of ASHRAE's specified comfort range for summers. We can clearly see that HAWADAAR achieves its goals by successfully adapting the tolerance range oscillating around the setpoint. Even at such high thermal loads, HAWADAAR is able to maintain a temperature tolerance of just ±0.4°C (see Figure 7 for 3PM onwards). Since the AC has the longest t_{safe}, these results represent the *worst case* of this best effort service.

5.2 Quantifying Energy Benefits of HAWADAAR

The are two main aspects of HAWADAAR's implementation that bring about its energy benefits. First, its standards-compliant thermal comfort resulting in aggressive duty cycling of single units.

Second, when available, its ability to prioritize low power devices. This section evaluates these two aspects of energy efficiency in multiple thermal settings.

Comparing energy savings in an experimental setting is not straight forward, as each day has its own parameter variations that affect how much energy should be spent to cool or warm a room. To draw logical conclusions, we still try to make approximate comparisons between days with similar average air temperature while minimizing the variance among internal thermal loads.

5.2.1 When Varying Tolerance Around Setpoint. We want to see if changing the tolerance level around a fixed setpoint affects the energy efficiency of the system. We use two tolerance ranges: ±0.5°C and ±1°C, thus satisfying category A (2°C) of ASHRAE's comfort requirement in terms of temperature deviation [4]. We keep the setpoint approximately at the middle of the comfort range: 22.5°C in winters and 24°C in summers, as depicted in Figure 8. For this part of the evaluation, the indoor sensor is placed on the wall opposite to the window sill, where the AC is installed (cf. Figure 3). We make the following key observations:

The nature of the heat transfer mechanism results in different thermal behaviors. The rise and fall of room temperature is gradual in winters (cf. Figures 8(a) and 8(b)) but rapid in summers (cf. Figures 8(c) and Figures 8(d)). In other words, the *period*—one complete on-off cycle—is larger for the heater in comparison with the AC as a consequence of the heat transfer mechanism. For the AC, cooling occurs through air convection. This can be short-lived after the AC is switched off as the air absorbs heat from the surrounding objects including external walls rapidly. For the heater, radiation is the physical mechanism of heat transfer, both in the form of visible and non-visible light, and effects all the surrounding objects not just the air. This process of heating is slower but longer lasting after the heater is switched off as the room air continues to absorb heat from heated objects maintaining a residual warm air temperature.

The nature of heat transfer has a contrasting impact on energy efficiency of Hawadaar *over different tolerance levels.* Table 2 summarizes the energy consumption and total on-time of devices per day. In winters, we can see that the on-time of the heater is reduced by approximately 3 hours when the tolerance increases from ±0.5°C to ±1°C. In contrast, in summers the AC on-time per day is similar for both tolerances. There are two reasons: First, the "slow start" of radiative heating makes it less efficient for frequent switching, as this increases the number of times the cold halogen element will have to be reheated (see magnifier in Figure 8(a)). Second, as described above, the ability of radiative heating to sustain the air temperature for longer durations given a larger tolerance range.

As an aside, duty cycling of the heater at ±0.5°C daily saves a further 11.6 kWh (40%) compared with continuous operation. Since ACs have onboard thermostats that already duty cycle, assuming similar setpoints and tolerance, such savings are not applicable.

5.2.2 When Exceeding Appropriate Temperature Settings. We now want to highlight the energy benefits of Hawadaar due to its ability to implement centralized policies, thus prohibiting excessive use of conditioning units. This refers back to our anecdotal observation in Section 1. Figure 9 shows how the AC on-time increases when the setpoint veers from recommended (24°C) to excessive (23, 22, and 21°C). We can see that even a single degree deviation from the recommended setpoint increases the on-time by 11%, which translates to ≈5 kWh per day. Similar observations, in the context of centralized HVAC, have been reported in Reference [25]. These results clearly reiterate the need for an Hawadaar-like solution to implement centralized policies in regions where the excessive and unhealthy use of conditioning units puts tremendous load on their stressed power grids.

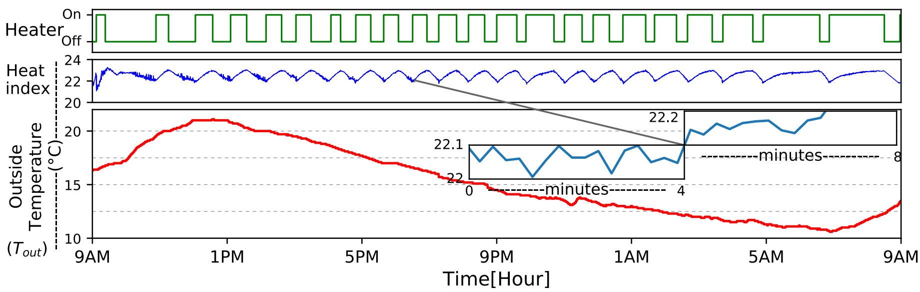

(a) Winter, January 25th (±0.5°C tolerance)

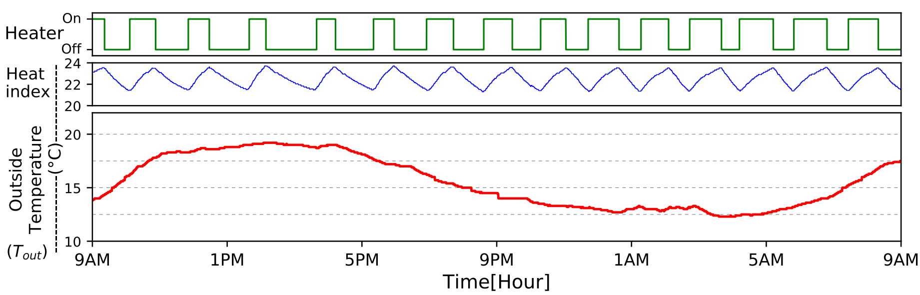

(b) Winter, January 27th (±1°C tolerance)

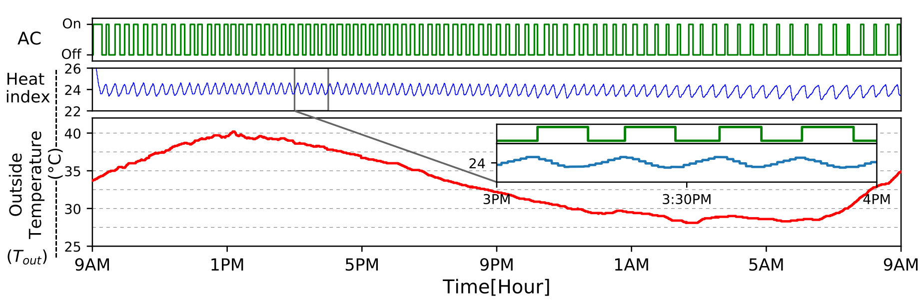

(c) Summer, May 8th (±0.5°C tolerance)

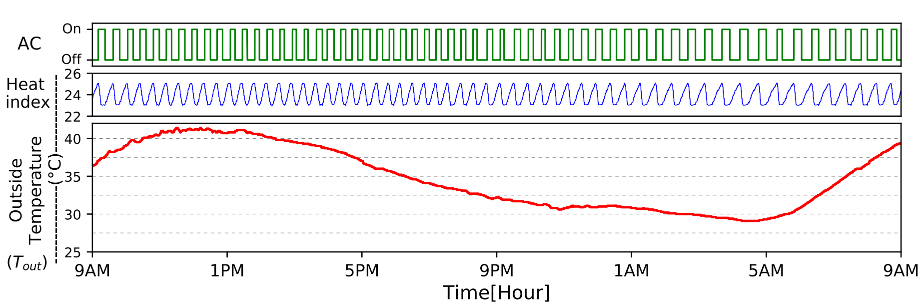

(d) Summer, May 10th (±1°C tolerance)

Fig. 8. Setpoint evaluation: Results are shown for two tolerance levels. The heat index reflects the room temperature. The thermal behavior in winters and summers is different due to the nature of heat transfer mechanisms: *convection* for AC and *radiation* for heater. The latter yields higher energy savings with wider temperature tolerance around the setpoint.

Table 2. Daily Energy Savings of HAWADAAR through Aggressive Duty Cycling of Single Units

Tolerance	On-time (hours/day)		Energy Consumption (kWh/day)		Energy Savings (kWh/day)	
	Heater	AC	Heater	AC	Heater	AC
±0.5	14.42(60%)	9.3(38%)	17.5	31.35	11.6	–
±1	11.26(47%)	9.27(38%)	13.6	31.15	15.5	–

Savings from AC are not applicable in this particular case as we expect the AC's onboard thermostat to achieve similar results for the same temperature settings.

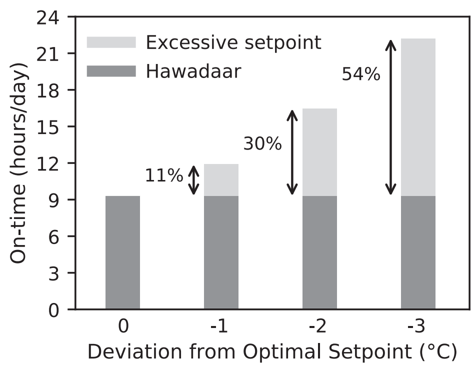

Fig. 9. Recommended vs. excessive: The standards-compliant HAWADAAR saves energy by prohibiting excessive setpoints.

5.2.3 By Interleaving Devices of Variable Energy Consumption. We now turn our focus on HAWADAAR's energy efficiency when thermal comfort is defined in terms of PMV. The idea is to evaluate the energy benefit of interleaving low-power ceiling fans. We measure this impact for ASHRAE's recommended range of PMV ($-0.5 < \text{PMV} < 0.5$) during two noticeably different weather conditions occurring at different times of the day in hot summers.

Day time results. Figure 10 shows results for maintaining PMV with (Figure 10(b)) and without (Figure 10(a)) a fan during hot day-time. We can observe that, under challenging weather conditions, the use of a fan has limited impact on the operational time of the AC, for two reasons. First, the ceiling fan pushes down the hot air that rises up after absorbing heat from surrounding objects, thus increasing air temperature. Second, the use of fan extends the duration for which PMV remains within specified range due to air circulation resulting in a further rise in air temperature. This increases the cooling load of AC during hot weather conditions.

Night time results. The impact of the fan is significantly pronounced under relatively moderate weather conditions at night, as can be seen in Figure 11. The prolonged use of ceiling fan helps maintain a desired PMV without the need for a high-power AC, resulting in ≈30% reduction in AC usage. These types of energy optimizations through aggregate usage of devices are highly unlikely through manual operation or device specific thermostat and definitely not possible at night with occupants asleep.

Overall, with PMV as comfort metric, we record a 15% reduction in AC usage per day, which translates into approximately 2.5kWh of energy savings after subtracting the power consumed by the fan.

6 POST ANALYSIS: ADAPTIVE VS OPTIMAL TWO-POSITION CONTROL

The adaptive two-position control strategy employed by HAWADAAR maintains a recommended setpoint while ensuring the operational safety of the thermal comfort devices. This is achieved by

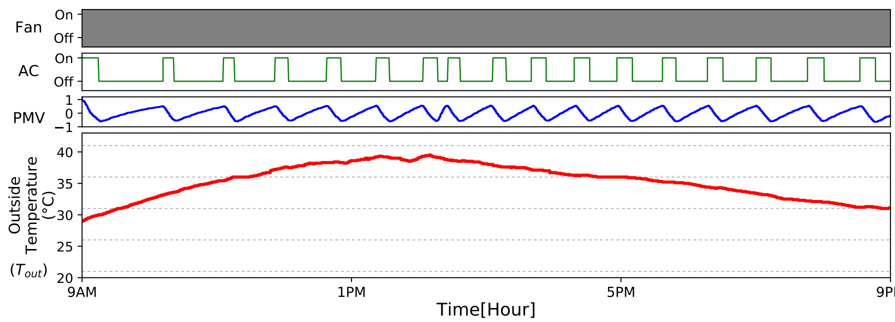

(a) PMV without fan (May 2nd)

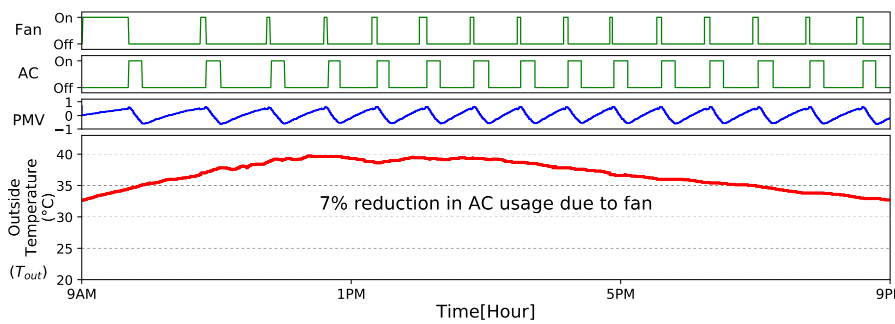

(b) PMV with fan (May 3rd)

Fig. 10. PMV evaluation (day): In hot conditions, interleaving a fan has limited impact on AC usage. The fan extends the duration for which PMV remains within specified range without using AC but results in a further rise in temperature, adding to the cooling load of AC.

always triggering control operations at the extrema of a given tolerance range around the setpoint, with a further constraint of always adjusting these extrema to meet the hard device constraints (t_{safe}), as discussed in Section 4.2. There is plenty of evidence to suggest that triggering the devices on or off at these extrema is not the most optimal in terms of energy efficiency.

6.1 Observations from Empirical Data

From our empirical data, the following observations are worth considering:

First, consider the on-time of heater during the interval 5am–9am in Figure 8(a). We can clearly see that the heater struggles to reach the control position T_{off}, i.e., the temperature at which the device is to be turned off. This results in an excessively prolonged on-time of the heater in cold weather conditions. This elicits a question: Could a different T_{off} lying within the same tolerance range, ±0.5°C in this case, be a better choice in terms of energy efficiency?

Second, consider the duty cycle of the AC at two different tolerance ranges, shown in Table 2. We can see that increasing the tolerance range from ±0.5°C to ±1°C hardly impacts the duty cycle of the AC. This suggests that the control points in our practical implementation are not necessarily optimal in terms of minimizing the duty cycle of the device.

Hawadaar's control algorithms greedily exploit a given tolerance range by triggering the control operation at the extrema. This approach is not necessarily the most energy efficient as indicated by the above observations. In the next section, based on the post analysis, we build a thermal model whose parameters are gleaned from the empirical data. The model helps us

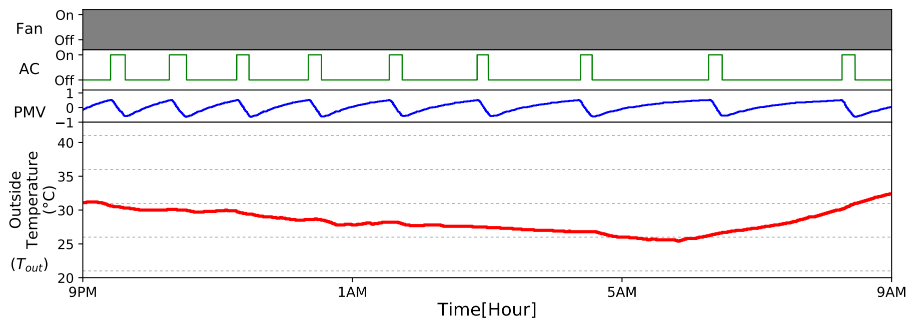

(a) PMV without fan (May 2nd/3rd)

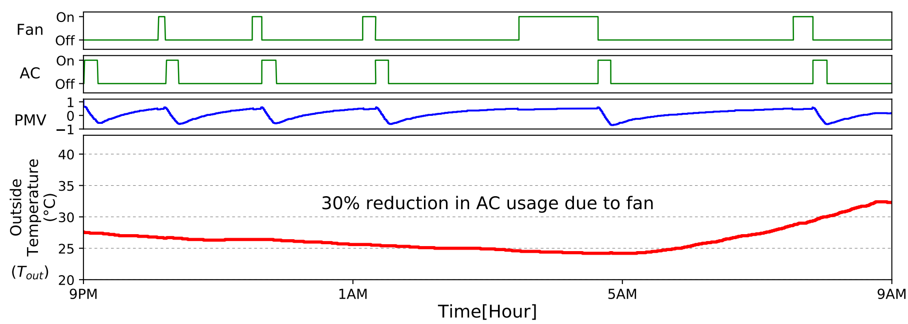

(b) PMV with fan (May 4th/5th)

Fig. 11. PMV evaluation (night): In moderate condition, the AC on-time is significantly reduced (≈30%) due to Hawadaar's ability to intelligently interleave the fan.

understand thermal capabilities of devices under varying weather conditions and find optimal control points in a given tolerance range provided complete information. These control points minimize the duty cycle of a device, thereby maximizing the energy efficiency. We would wish to evaluate how far is Hawadaar from such a theoretically optimal control strategy.

6.2 Building a Thermal Model

The heating/cooling operation of a device can be divided into multiple stages in terms of how it impacts the ambient temperature.

When a device is initially turned on, it gradually starts conditioning to capacity after some delay. For a heater, this delay corresponds to the time taken by the halogen elements to heat up. For an AC, this delay is the time elapsed in minimizing the pressure differential of the gases on both sides of the compressor. This results in a *lag* stage, as shown in Figure 12(a), before the ambient temperature is affected. After the lag stage is the *exponential* stage, where the ambient temperature changes at a faster rate. Next to the exponential stage is the *saturation* stage when running the device further hardly impacts the ambient temperature (see also time period 5am–9am in Figure 8(a)). A similar reasoning can be used to describe the changes in ambient temperature when a device is turned off. For instance, the lag stage representing the delay when the heater/AC is turned off and when it actually stops heating/cooling the room, due to the dynamics of the halogen elements and heat pumps that do not halt their conditioning operation instantly. The subsequent two stages can also be inferred from the Newton's law of cooling.

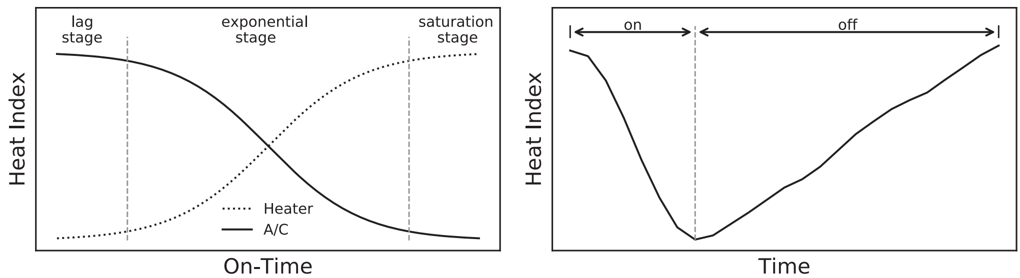

(a) Heating/cooling stages (b) One complete period of AC's operation

Fig. 12. Describing thermal behavior of heater and AC: (a) a slow-start (lag stage) is followed by changes in ambient temperature at a faster rate (exponential stage) until the device hardly impacts the ambient temperature (saturation stage); (b) a randomly selected on-off period of the AC mimicking the described behavior.

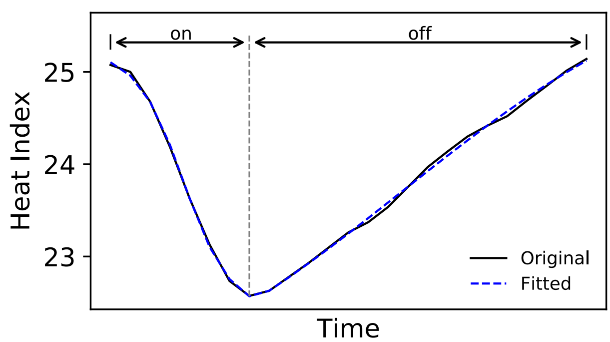

Fig. 13. Curve fitting of the thermal model on a single period of AC operation.

Figure 12(b) depicts a randomly selected on-off period of the AC from our empirical data. We can clearly see the different stages of thermal behavior described in Figure 12(a). Identifying these different stages is vital in understanding the impact of a device's operation on the ambient temperature, thereby finding the optimal control points to minimize the duty cycle. We represent this type of thermal behavior using a generalized logistic function:

$$Y(t) = D + \frac{A}{1 + e^{-B(t-C)}}, \tag{6}$$

where A is a linear combination of the lower and upper asymptote; B is the growth rate; C corresponds to the starting time; and D is the lower asymptote.

From the empirical data collected over a 7-month deployment, we observe that each period consists of two sigmoidal curves (one for the *on* and the other for the *off* duration) with varying A, B, C, and D parameters. As the duration of the period changes due to varying thermal load, these parameters vary for every period. We derive these model parameters from the empirical data *separately* for each on and off duration of all periods. We employed a non-linear least-square (NLS) problem setup for curve fitting, while using a Levenberg-Marquardt algorithm to solve this NLS problem. We also restricted the parameter learning to remain within the tolerance range of respective experiments. This fine-grained parameter learning results in high fitting accuracy as can be seen in Figure 13 for a randomly selected conditioning period of the AC. Over a day, this results in a mean residual error of only 0.032 and 0.029 and standard deviation of 0.039 and 0.036

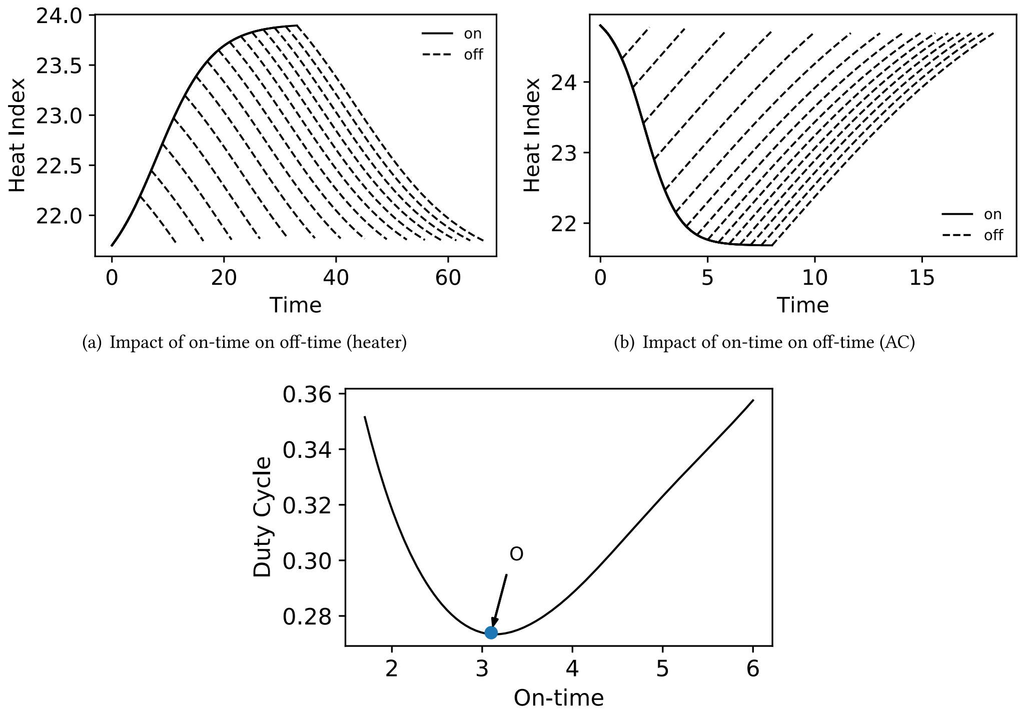

(a) Impact of on-time on off-time (heater)

(b) Impact of on-time on off-time (AC)

(c) Duty cycle vs on-time

Fig. 14. Impact of varying on-time upon off-time: (a, b) increasing the on-time initially has a larger impact upon the off-time but subsequently the impact decreases for both heater and AC; (c) correspondingly, the duty cycle is a convex curve with a minimum at point O.

for the on and off durations, respectively. Similarly, our results achieve a convergence tolerance of magnitude 1.49e-08.

We next discuss how to use this model to develop an optimization problem with the objective to minimize the duty cycle of a device's operation.

6.3 Defining the Optimization Problem

Define *tolerances*, such as ±0.5°C or ±1°C, about a setpoint T_s. The upper extrema U_b is at $T_s + \frac{\text{tolerance}}{2}$, while the lower extrema L_b is at $T_s - \frac{\text{tolerance}}{2}$. Starting with a temperature L_b, as shown in Figure 14(a) for the heater, we study the cyclical change in the temperature for varying on and off times. Initially, increasing the on-time has a larger impact on the off-time; however, later this impact gradually decreases. A similar trend can be found in Figure 14(b) for the AC when starting with a temperature U_b. Thus, we can conclude that increasing the on-time does not proportionally increase the off-time. Our goal is to find the minimal on-time that would result in the largest off-time for a given period, thus minimizing the duty cycle of a device for maximum energy savings. The on-time for which the duty cycle defined below is minimum will precisely represent the optimal control point to turn off the device to minimize the duty cycle:

$$\text{duty cycle} = \frac{\text{on-time}}{\text{period}}. \tag{7}$$

The same relationship is shown in Figure 14(c). Initially, the duty cycle decreases up to a certain point O and then increases afterwards, conveniently representing a convex curve. Hence, O is the optimal control point for turning off the device during a given period in terms of minimizing the energy consumption of the device. *An optimal approach in terms of energy saving would be one that can find* O. However, the point O depends upon the exact thermal path traced during a period and requires *a priori* information regarding its cyclical variation, which may change depending upon a number of external factors (external weather, number of occupants, etc.). Thus, optimal energy savings may only be evaluated based on a post analysis of the empirical data. Furthermore, we need to find O whilst satisfying the soft thermal and hard device constraints described in Section 3. A formal description of this optimization problem is given below:

$$
\left.
\begin{array}{rl}
 & \text{Optimization Model} \\
\text{Notations} & \\
q_1(t) & \text{temperature during the on-time, } 0 \le t \le t_1 \\
q_2(t) & \text{temperature during the off-time, } t_1 \le t \le p \\
T_s & \text{average temperature to be maintained} \\
L_b & \text{lower extrema} \\
U_b & \text{upper extrema} \\
 & \\
\text{Decision Variables} & \\
t_1 & \text{on-time} \\
t_2 & \text{off-time} \\
p & \text{period} \\
t_o & \text{is the epoch when the AC is turned off} \\
\text{Minimize} & \dfrac{t_1}{p} \\
\text{Subject to} & \\
 & q_1(t) = \dfrac{a_1}{1 + e^{-B_1(t-C_1)}} + d_1 \quad 0 \le t \le t_o \\
 & q_2(t) = \dfrac{a_2}{1 + e^{-B_2(t-C_2)}} + d_2 \quad t_o \le t \le p \\
 & q_1(t_o) = q_2(t_o) \\
 & T_s = \dfrac{\int_0^{t_o} q_1(t)dt + \int_{t_o}^{p} q_2(t)dt}{p} \\
 & L_b \le q_1(t), q_2(t) \le U_b \\
 & t_2 \ge t_{safe}
\end{array}
\right\}
\tag{8}
$$

6.4 Results

We now apply this model-based control approach to find the optimal control points for turning the device on (T_{on}) and off (T_{off}) for the data depicted in Figure 8, and we evaluate whether T_{off} and T_{on} still lie on the extrema of the tolerance range. If this is true, then Hawadaar's control approach would have achieved the minimum duty cycle of a device for a given tolerance range. Otherwise, we would like to assess how far is Hawadaar from an approach that is aware of the optimal control points to minimize the duty cycle. Figure 15 shows sample path results for winters and summers

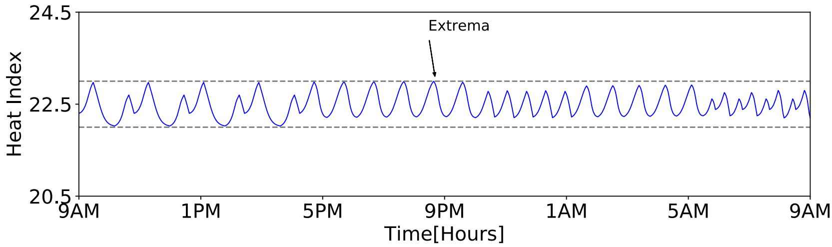

(a) Winter, January 25th (±0.5°C tolerance)

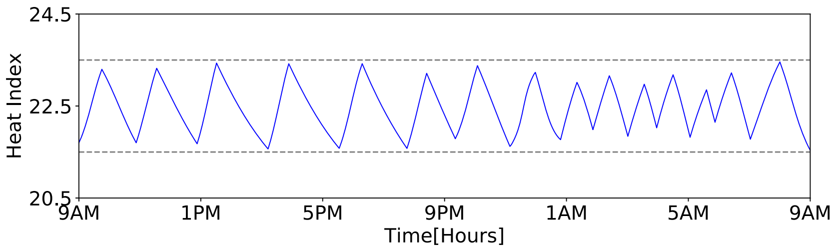

(b) Winter, January 27th (±1°C tolerance)

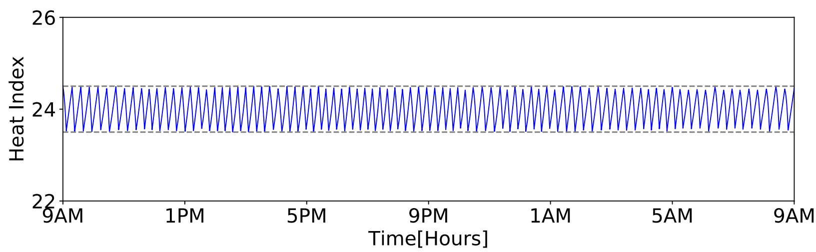

(c) Summer, May 8th data (±0.5°C tolerance)

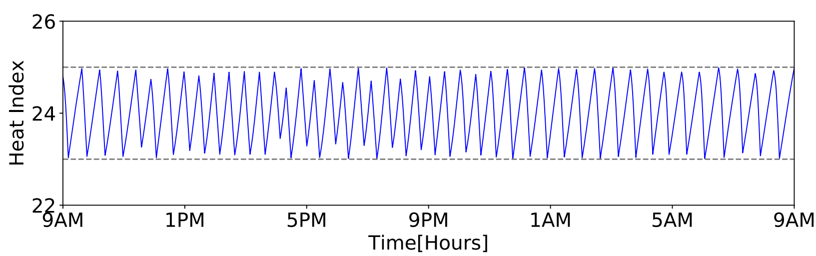

(d) Summer, May 10th data (±1°C tolerance)

Fig. 15. Sample path traced by model-based control: the optimal control points do not necessarily lie on the extrema of the tolerance range as used by HAWADAAR to trigger its control operations.

using the two tolerance ranges, ±0.5°C and ±1°C. We can clearly see that these optimal control points for a given tolerance range do not necessarily lie on the extrema. The resulting reduction in the duty cycles is summarized in Table 3, with a model-based approach achieving a 9% reduction, on average.

The difference is understandably bigger when the tolerance range gets wider as shown in the last two columns of Table 3 for both heater and AC. A wider tolerance range increases the search

Table 3. Model-Based vs. Practical Implementation of Two-position Control Strategy

	On-time (hours/day)					
	Model-based		HAWADAAR		Difference	
Tolerance	**Heater**	**AC**	**Heater**	**AC**	**Heater**	**AC**
±0.5	13.83	8.78	14.42	9.30	4%	6%
±1	9.9	8.18	11.26	9.27	14%	13%

A model-based approach can reduce the duty cycle by 9%, on average.

space as well as the chances of finding a better control point below the extrema, and hence the corresponding reduction in the duty cycle. However, such an optimal approach requires complete information regarding the cyclical variations in ambient temperature, which depends on a wide range of factors such as room insulation, devices under study, number of occupants, time of the day, external temperature and humidity, and so on. These further savings thus represent an upper bound to the possible savings that might be accrued with a model predictive approach instead of HAWADAAR's simple, adaptive two-position control.

7 ESTIMATING COUNTRYWIDE BENEFITS

To illustrate the energy benefits of HAWADAAR on a countrywide scale, we use the results from Section 5, together with estimates of annual electricity consumption on air conditioners and electric heaters in residential and commercial buildings. We are interested in extrapolating energy savings for an emerging economy that is populous, with a sizeable middle and upper middle class in urban areas living in buildings that can be retrofitted with HAWADAAR, and faces extreme climate with significant temperatures variations in summer and winter months.

The main challenge we encounter in this exercise is the lack of electricity consumption estimates in buildings (residential or commercial), at an aggregate level or by end use, in such a context. To address this, we employ an international benchmarking approach that uses plausible inputs from U.S. electricity consumption surveys, to arrive at conservative estimates of annual electricity consumption in residential and commercial buildings on air conditioners and electric heaters in our typical economy. The inputs and assumptions underlying our projected savings are discussed below.

7.1 Inputs and Assumptions

The first input is the annual electricity consumption by end use in buildings in the U.S. as recorded in the *Residential Energy Consumption Survey* (RECS) [15] and *Commercial Building Energy Consumption Survey* (CBECS) [6]. The RECS provides data on annual electricity usage in kWh for residences, while the CBECS provides data on electricity usage in kWh and floor space in square feet for commercial buildings in different climate regions of the U.S. by end use, such as air conditioning and space heating, amongst other uses. The simplifying assumption we make is that U.S. electricity consumption data is representative of electricity usage in emerging economies. There are undoubtedly variations in building materials, appliance efficiency, and engineering systems in the U.S. relative to any developing country, and we should expect the U.S. consumption to be lower when accounting for better technologies for cooling and heating. This would lead us to start out with more conservative estimates of consumption in emerging economies, and thus more conservative estimates of savings. To ensure that consumption patterns due to seasonal variations are similar to our context, we take care to focus only on buildings in hot-humid, mixed-humid, and mixed-dry/hot-dry building regions. These climate regions were created by the Building America

Table 4. Inputs for Estimating Total Annual Electricity Saving

	Average	Hot Humid	Mixed-Dry/Hot-Dry	Mixed-Humid
Residential AC ($kWh/household$)	2,442.00	4,077.00	1,899.00	1,777.00
Commercial AC (kWh/ft^2)	3.10	5.34	2.10	2.52
Commercial Heating (kWh/ft^2)	0.34	0.12	0.15	0.52
Population				
Number of households	10 million			
Commercial Floor Space	1000 million square feet			
Penetration Rate	0–50%			

program and are meant to capture the differences in climate and building types in different parts of the country. These regions include cities such as Houston, Dallas, Pheonix, Memphis, and Atlanta, which have climate similar to cities such as Mumbai, New Delhi, Karachi, Lahore, Tehran, Dhaka, Beijing, and Cairo.

The second input is the stock of residential and commercial buildings that can benefit from Hawadaar. We calculate energy savings for a large emerging economy, with significant urbanization. We assume a total population of 250 million persons, urbanization rates of 40%, implying an urban population of 100 million persons. Assuming that half of the urban population lives in buildings where Hawadaar can be used, and average household consists of five persons, we have a beneficiary population of 10 million households or residential units. To arrive at estimates of commercial floor space, we use plausible numbers of the density of commercial floor space per capita of various developing countries available from recent research [19], assuming that the stock of commercial floor space is 1,000 million square feet.

The third input is the penetration rate of Hawadaar in our beneficiary population. In the base case, we assume that the average penetration rates of Hawadaar for air conditioners is 20%. For electric heaters, we assume 20% penetration in commercial buildings and 0% in residential units, as we expect natural gas-based space heating to be predominant in households due to its cheaper cost; currently not supported by Hawadaar.

Finally, based on our empirical results in Section 5, we normalize energy savings emanating from both the key features of Hawadaar: (i) aggressive duty cycling of single units and (ii) interleaving of low-power devices. In the first case, we set energy savings from AC to 30% (cf. Section 5.2.2): This is the *optimistic* case, where we assume an average deviation of -2°C from the optimal setpoint based on our anecdotal observation. Hawadaar can preempt such excessive settings to claim its energy savings because it is standards-compliant. In the second case, we assume the energy savings from AC to 15% (cf. Section 5.2.3). This is the *pessimistic* case, where we disregard our anecdotal observation and only consider savings from Hawadaar's ability to minimize energy consumption through interleaving low-power devices when using PMV as a comfort metric. We strongly believe that this is the bare minimum benefit of deploying Hawadaar. Finally, in both cases the energy savings from electric heater are set to 40% (cf. Section 5.2.1), assuming no onboard temperature control units and that only commercial buildings use this type of heating.

Our inputs to calculate projected savings are summarized in Table 4. Total savings are found by multiplying the per unit air conditioner and electric heating consumption for different climate zones with savings above, assuming a penetration rate of 20%. We also calculate savings for different penetration rates. With these inputs, we calculate very conservative estimates of total savings (cf. Figure 16), which can be considered a lower bound on actual savings that may be realized from implementing Hawadaar.

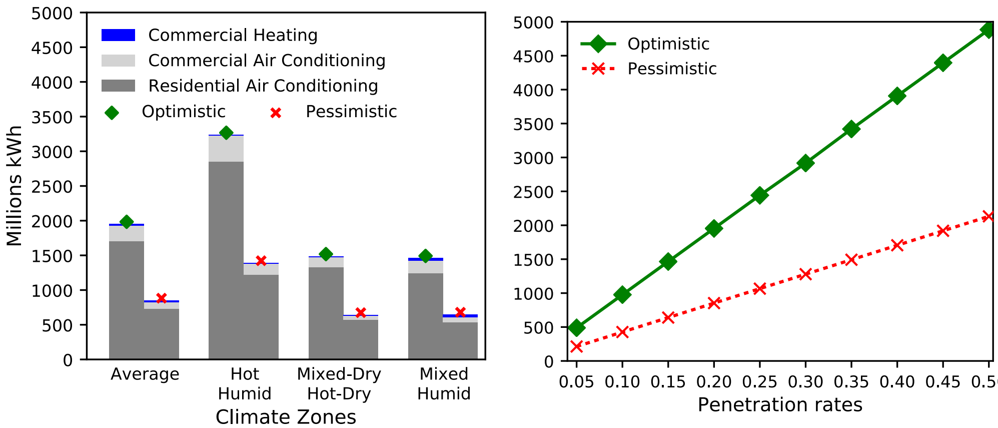

(a) By climate zones at fixed HAWADAAR penetration rate of 20%.

(b) By varying market penetration.

Fig. 16. Countrywide estimates of annual energy savings.

7.2 Savings and Environmental Impact

Figure 16(a) shows the countrywide estimate of energy savings by different climate zones at a market penetration of 20%. In the optimistic scenario, savings range from 1,462 million kWh to 3,237 million kWh, and in the pessimistic scenario, savings range from 645 million kWh to 1,393 million kWh. Savings are higher for cooling in humid zones, given the higher consumption per household and per unit area. Similarly, heating consumption and thus savings are higher in climate zones with greater need for heating during the year. Given our assumptions about the population size, residential air conditioning accounts for ≈85%, while commercial air conditioning and heating together account for 15% of total estimated savings. In Figure 16(b), we display the effect of changing the penetration rates on savings for the buildings located in average climate zone. Total savings rise proportionally with penetration rates, as we assume a uniform consumption rate in our population.

To understand the magnitude of total energy savings, we can also express them as a fraction of per capita consumption in the residential and commercial sectors. The average per capita electricity consumption in lower middle income countries ranges from 500 to 1,000 kwh per capita out of which approximately 30% is residential and commercial electricity consumption [8]. Using the midpoint consumption of 750 kwh per capita, we can conclude that residential and commercial sectors account for 255 kwh per capita of electricity consumption. Then, we can express our total savings in per capita terms and find the percent saved. The estimates in Figure 16(a) imply that between 1.15% to 5.76% of electricity consumption in these sectors can be saved with penetration rates of *just* 20%. As we increase the penetration rates from 5% to 50%, the percentage savings range from 0.87% to 8.68% in the optimistic scenario and from 0.38% to 3.79% in the pessimistic scenario. Overall, these projections illustrate that even with very conservative assumptions, HAWADAAR can have an economically meaningful impact on consumption at a macro level.

Furthermore, the projected energy savings directly translate into reduction in carbon emissions, for a given carbon dioxide to kWh emission factor. Therefore, we should expect countrywide emissions per capita due to electricity generation to fall by the same percentage as the reduction in electricity consumption per capita. Thus, in emerging economies where fossil fuels are used heavily in electricity generation, and carbon emissions are large in magnitude, the total value of reduction in emissions will be larger still. We therefore believe in a much wider impact of HAWADAAR in the campaign against global warming, as its implementation could possibly accredit emerging countries for a major role in Kyoto Protocol [27] by accepting binding targets of reduction in carbon emissions, without compromising on the quality of life improvement for their citizens.

8 RELATED WORK AND OUTLOOK

Related Work: Efficient operation of HVAC has been at the forefront of existing literature on *energy conservation* [3, 7, 9, 24, 31–33] and *thermal comfort* [16, 18, 30, 34] in buildings. Energy conservation is typically achieved by incorporating occupancy patterns [7, 24, 33] and thermal load predictions [5, 9] into HVAC operation schedules, or by exercising more fine-grained control, such as room-level air flow control [31] and stage selection in a multi-stage HVAC [32]. Aswani et al. [5] employ similar predictive techniques for the energy efficient operation of a room-level AC only, thereby addressing a subset of problems considered in this article. Studies focusing on improving thermal comfort of centralized HVAC try to find improved setpoints based on occupants' feedback [16, 34], or by augmenting HVAC with personal devices to create micro thermal-zones around a user for highly personalized thermal comfort [2, 18, 30]. These personal comfort systems nonetheless rely on HVAC to first achieve a building wide setpoint; this facility is not available in our setup.

Our goal to achieve thermal comfort at low energy budget is aligned with existing literature but the nature of challenges we face is inherently different. For example, the type of buildings, the extreme weather, as well as thermal characteristics and location of conditioning units entail us to build more aggressive control strategies, such as the ones employed by Hawadaar. This article thus primarily focuses on developing and evaluating these centralized control strategies. While the current implementation of Hawadaar is oblivious to occupancy prediction (estimating the number of occupants for determining internal thermal load) due to its reactive control strategy, occupancy detection (if and when the room is occupied) is an orthogonal but well-researched problem outside the current scope of this article. Thus, we do not foresee any inherent challenges in the seamless incorporation of existing occupancy detection solutions in Hawadaar, to refine its operation schedules and further reduce energy consumption.

The work that comes closest to our idea of interleaving AC and ceiling fan is the collaboration of NEST with a smart ceiling fan company [21]. The key idea is to adjust fan's speed as temperatures rise, allowing to increase thermostat setpoint of a centralized HVAC while still feeling just as cool. However, we observed that for our range of operating temperature and suboptimal building insulation, simultaneous use of AC and fan results in higher average temperature as the fan forces hot air down.

Outlook: Given a certain thermal objective, we have shown that Hawadaar can operate within strict comfort bounds by either utilizing a single device or interleaving multiple devices supporting different modes of air conditioning. Going forward, our primary question is to find out how comfortable is Hawadaar for occupants? This requires an extensive user experience study. We are also interested in expanding Hawadaar's control to include more devices, such as ventilation fans and evaporative coolers, which are widespread among the middle-income groups in emerging countries. Similar to ceiling fans, these devices are low-power and thereby provide greater saving opportunities through interleaving. For example, a bidirectional ventilation fan offers us the ability to allow fresh air (and more oxygen) in or hot air out, and an evaporative cooler can help control air temperature and humidity in hot and dry weathers. Finally, we would also like to evaluate our system in larger rooms—labs, classrooms, halls, and so on—which are fitted with multiple air conditioning units. In addition to greater internal thermal load, such spaces also pose further control challenges; for example, how to orchestrate the operation schedule of multiple ACs that reduces peak load without violating comfort bounds?

9 CONCLUSIONS

This article serves as a proof of concept for the fundamental components of an unconventional, inverted HVAC architecture for older buildings in emerging countries. As an alternative to modern HVAC, we proposed IoT-based retrofits to reinforce legacy air conditioning units for policy driven actuation. Hawadaar is a practical realization of this proposal, demonstrating its efficacy in achieving a *high* level of thermal comfort at *low* energy budget. Our post analysis shows that the efficiency of Hawadaar, in terms of energy consumption, is not far from a theoretically optimal approach that has the the optimal control points for minimizing the duty cycle. Our post analysis shows that Hawadaar's adaptive two-position control is not far from a theoretically optimal approach based on *a priori* information that precisely knows the optimal control points to minimize consumption. Our empirical evaluations, when plausibly scaled to countrywide estimates, highlight the worthwhile impact of Hawadaar, providing energy savings that directly translate into reduced carbon emissions in countries that rely heavily on burning fossil fuels for electricity generation. We see a greater value in pursuing this work further not only to widen its impact, such as by including more device types, but also to improve its implementation and algorithmic aspects.

REFERENCES

[1] Amana. 2018. Room air conditioner & heat pump: Use and care manual. Retrieved from www.amana.com.

[2] Michael P. Andersen, Gabe Fierro, Sam Kumar, Michael Chen, Leonard Truong, Joyce Kim, Edward A. Arens, Hui Zhang, Paul Raftery, and David E. Culler. 2015. Well-connected microzones for increased building efficiency and occupant comfort. In *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*. ACM, 121–122.

[3] Omid Ardakanian, Arka Bhattacharya, and David Culler. 2016. Non-intrusive techniques for establishing occupancy related energy savings in commercial buildings. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys'16)*. ACM, New York, NY, 21–30. DOI: https://doi.org/10.1145/2993422.2993574

[4] ASHRAE. 2010. Thermal environmental conditions for human occupancy, standard 55. Retrieved from http://arco-hvac.ir/wpcontent/uploads/2015/11/ASHRAE-55-2010.pdf.

[5] Anil Aswani, Neal Master, Jay Taneja, David Culler, and Claire Tomlin. 2012. Reducing transient and steady-state electricity consumption in HVAC using learning-based model-predictive control. *Proceedings of the IEEE* 100, 1 (Jan. 2012), 240–253. https://ieeexplore.ieee.org/document/5985456.

[6] CBECS 2012 (Tables B27 and E5). 2012. Commercial buildings energy consumption survey. Retrieved from https://www.eia.gov/consumption/commercial/data/2012.

[7] Bharathan Balaji, Jian Xu, Anthony Nwokafor, Rajesh Gupta, and Yuvraj Agarwal. 2013. Sentinel: Occupancy-based HVAC actuation using existing WiFi infrastructure within commercial buildings. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems (SenSys'13)*. ACM, New York, NY, Article 17, 14 pages. DOI: https://doi.org/10.1145/2517351.2517370.

[8] World Bank. 2016. World development indicators. Retrieved from http://data.worldbank.org/data-catalog/world-development-indicators.

[9] Alex Beltran and Alberto E. Cerpa. 2014. Optimal HVAC building control with occupancy prediction. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings (BuildSys'14)*. ACM, New York, NY, 168–171. DOI: https://doi.org/10.1145/2674061.2674072

[10] Weather Prediction Center. [n.d.]. Retrieved from http://www.wpc.ncep.noaa.gov.

[11] Stan Cox. 2010. *Losing Our Cool: Uncomfortable Truths About Our Air-Conditioned World (and Finding New Ways to Get Through the Summer)*. The New Press. Retrieved from http://thenewpress.com/books/losing-our-cool.

[12] Stan Cox. 2012. Cooling a warming planet: A global air conditioning surge. *Yale Environ.* 360, 10. Retrieved from http://e360.yale.edu.

[13] Lucas W. Davis and Paul J. Gertler. 2015. Contribution of air conditioning adoption to future energy use under global warming. *Proc. Natl. Acad. Sci. U.S.A.* 112, 19 (2015), 5962–5967.

[14] United Nations Population Devision. 2015. Revision of world population prospects. Retrieved from https://esa.un.org/unpd/wpp/.

[15] RECS (Table C E41). 2009. Residential energy consumption survey. Retrieved from https://www.eia.gov/consumption/residential/data/2009/.

[16] Varick L. Erickson and Alberto E. Cerpa. 2012. Thermovote: Participatory sensing for efficient building HVAC conditioning. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys'12)*. ACM, New York, NY, 9–16. DOI: https://doi.org/10.1145/2422531.2422534

[17] Frigidaire. 2018. All about the use and care of your room air conditioner: User manual. Retrieved from www.frigidaire.com.

[18] Peter Xiang Gao and S. Keshav. 2013. Optimal personal comfort management using SPOT+. In *Proceedings of the 5th ACM Workshop on Embedded Systems for Energy-Efficient Buildings (BuildSys'13)*. ACM, New York, NY, Article 22, 8 pages. DOI: https://doi.org/10.1145/2528282.2528297

[19] Burak Güneralp, Yuyu Zhou, Diana Ürge-Vorsatz, Mukesh Gupta, Sha Yu, Pralit L. Patel, Michail Fragkias, Xiaoma Li, and Karen C. Seto. 2017. Global scenarios of urban density and its impacts on building energy use through 2050. *Proc. Natl. Acad. Sci. U.S.A.* 114, 34, 8945–8950.

[20] Khadija Hafeez, Yasra Chandio, Abu Bakar, Ayesha Ali, Affan A. Syed, Tariq M. Jadoon, and Muhammad Hamad Alizai. 2017. Inverting HVAC for energy efficient thermal comfort in populous emerging countries. In *Proceedings of the 4th ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys'17)*. ACM, New York, NY, Article 18, 10 pages. DOI: https://doi.org/10.1145/3137133.3137137

[21] Haiku-fans. 2018. Retrieved from https://www.haikuhome.com/.

[22] International Energy Angency (IEA). 2013. Policy pathways: Modernising building energy codes. https://www.iea.org/publications/freepublications/publication/PolicyPathwaysModernisingBuildingEnergyCodes.pdf.

[23] Morna Isaac and Detlef P. van Vuuren. 2009. Modeling global residential sector energy demand for heating and air conditioning in the context of climate change. *Energy Policy* 37, 2 (2009), 507–521. Retrieved from https://EconPapers.repec.org/RePEc:eee:enepol:v:37:y:2009:i:2:p:507-521.

[24] Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse. 2010. The smart thermostat: Using occupancy sensors to save energy in homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys'10)*. ACM, New York, NY, 211–224. DOI: https://doi.org/10.1145/1869983.1870005

[25] M. M. Manning, M. C. Swinton, F. Szadkowski, J. Gusdorf, and K. Ruest. 2005. The effects of thermostat setting on seasonal energy consumption at the CCHT research facility. Canada Center for Housing Technology. https://nparc.nrc-cnrc.gc.ca/eng/view/object/?id=344f31a2-8d67-4c5a-84bd-488fe7a18f59.

[26] Michael A. McNeil and Virginie E. Letschert. 2008. Future air conditioning energy consumption in developing countries and what can be done about it: The potential of efficiency in the residential sector. Retrieved from http://escholarship.org/uc/item/64f9r6wr.

[27] United Nations. 1998. Kyoto protocol to the United Nations framework convention on climate change. Retrieved from http://unfccc.int/resource/docs/convkp/kpeng.pdf.

[28] Karsten Neuhoff, Hermann Amecke, Aleksandra Novikova, and Kateryna Stelmakh. 2011. Thermal efficiency retrofit of residential buildings: The German experience. In *CPI Report, Climate Policy Initiative*. Retrieved from http://hdl.handle.net/10419/65868.

[29] F. Oldewurtel, A. Parisio, C. N. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and K. Wirth. 2010. Energy efficient building climate control using stochastic model predictive control and weather predictions. In *Proceedings of the American Control Conference*. IEEE, 5100–5105.

[30] Alimohammad Rabbani and S. Keshav. 2016. The SPOT* personal thermal comfort system. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys'16)*. ACM, New York, NY, 75–84. DOI: https://doi.org/10.1145/2993422.2993578

[31] Virginia Smith, Tamim Sookoor, and Kamin Whitehouse. 2012. Modeling building thermal response to HVAC zoning. *SIGBED Rev.* 9, 3 (July 2012), 39–45. DOI: https://doi.org/10.1145/2367580.2367587

[32] Tamim Sookoor and Kamin Whitehouse. 2013. RoomZoner: Occupancy-based room-level zoning of a centralized HVAC system. In *Proceedings of the ACM/IEEE 4th International Conference on Cyber-Physical Systems (ICCPS'13)*. ACM, New York, NY, 209–218. DOI: https://doi.org/10.1145/2502524.2502553

[33] Amee Trivedi, Jeremy Gummeson, David Irwin, Deepak Ganesan, and Prashant Shenoy. 2017. iSchedule: Campus-scale HVAC scheduling via mobile WiFi monitoring. In *Proceedings of the 8th International Conference on Future Energy Systems (e-Energy'17)*. ACM, New York, NY, 132–142. DOI: https://doi.org/10.1145/3077839.3077846

[34] Daniel A. Winkler, Alex Beltran, Niloufar P. Esfahani, Paul P. Maglio, and Alberto E. Cerpa. 2016. FORCES: Feedback and control for occupants to refine comfort and energy savings. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp'16)*. ACM, New York, NY, 1188–1199. DOI: https://doi.org/10.1145/2971648.2971700

[35] Wireless Female Zwave Power Plug. 2018. Retrieved from https://qianpeng.en.alibaba.com/.

[36] Zwave Temperature Humidity Sensor. 2018. Retrieved from https://heimansmart.en.alibaba.com/.

Received January 2018; revised May 2018; accepted May 2018