Understanding and Reducing HVAC Power Consumption Post Evacuation Events in Commercial Buildings

Iresha Pasquel Mohottige, Hassan Habibi Gharakheili, Arun Vishwanath, Salil S. Kanhere, Vijay Sivaraman

***Abstract*—Buildings are required to follow standard operational procedures during emergency evacuation. In addition to people evacuating the building, one of the recommended steps during a fire evacuation is to shut down the air handling units (AHUs) of the heating, ventilation and air conditioning (HVAC) system to prevent smoke from spreading in the building via the air ducts. Shutting down the AHU will inevitably cut off cooling, resulting in internal temperatures rising steeply particularly on hot days. This phenomenon imposes considerable power demand on the HVAC to rapidly cool the building down during reoccupation. In this paper, we study the energy implications of post evacuation scenarios. Our contributions are three-fold: (1) We quantify power excursion caused in 43 evacuation events across 14 buildings of a university campus using a data-driven building thermal model. We show evacuations during summer season can result in power consumption up to 150% above the power demand threshold. (2) We develop a method to reschedule planned evacuations in order to eliminate the power excursions while adhering to building evacuation standards. (3) We develop a formal optimization framework to minimize the energy costs during planned and emergency evacuations without compromising desired thermal comfort temperatures by intelligently cooling the building post evacuation. This is the first study to understand and reduce the HVAC power consumption associated with building evacuation events.**

***Index Terms*—HVAC power consumption, evacuation, WiFi occupancy, thermal model**

I. Introduction

MOST fatalities during a building fire emergency are not due to the fire but rather due to suffocation attributed to smoke spreading through the building [15]. While building heating, ventilation and air conditioning (HVAC) is necessary to provide adequate climate control for occupants, an unintended consequence of keeping the ventilation system (*i.e.,* fans) operational during a fire is smoke spreading rapidly to the non-affected sections of the building via the air handling ducts. To mitigate this disastrous outcome, the HVAC is automatically switched to a fire-safe mode, either via a dedicated smoke control system, which either turns of the AHU or the entire HVAC system [22]. Either way, altering the operation of the HVAC can cause the internal zone temperatures to deviate significantly from the desired comfortable (*i.e.,* set-point) temperatures. The duration of the shut down and the current outside air temperature may further exacerbate the situation. Deactivating the HVAC system (due to a fire evacuation) during noon hours on a hot summer day, can raise the indoor temperature to be equal to the ambient outdoor temperature in a few tens of minutes [10], [18].

Once an all-clear is given post evacuation, building reoccupation begins. During this time, the HVAC system typically runs at full load to cool the building down rapidly and bring the zone temperatures to more comfortable levels as determined by the configured set-points. A significant power draw during this event can adversely impact the electricity bills paid by the building owners. Electricity suppliers maintain reserve generation capacity as stand-by to respond to transient power demands. However, they charge a significant premium for such service. Note that up to 70% of the building energy demand can be attributed to HVAC alone [3], [30], [35].

More important than drills (planned) are emergency evacuations (unplanned) that can occur more frequently than expected, especially in older buildings [39], with no control over their occurrence time. Thus, power excursions post evacuations may result in the premise owners being penalized for exceeding the peak demand threshold that is set out in their contract. To the best of our knowledge, no prior work quantifies the impact of evacuation events on the cost of energy. The peak demand charge can constitute between 50% and 70% of the monthly electricity bill [14], [16]. At the same time, occupant thermal comfort is a key concern of the building managers. Therefore, it is challenging to achieve a balance between the energy costs and the occupant thermal comfort during post evacuation scenarios. While numerous prior works exist like [35] on quantifying and managing the daily cost of HVAC systems, our work is first to analyze the energy implications of atypical events such as evacuations.

For this work, we obtain WiFi session traces data (for computing building occupancy), and a list of planned and emergency evacuation events (detected by an automatic method which we developed in our prior study [23]) in 14 buildings of a large university campus during a period of 180 days. For our **first** contribution, using real-data we show that excursions occur in cooling power demand of buildings

I. Pasquel Mohottige, H. Habibi Gharakheili, and V. Sivaraman are with the School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, NSW 2052, Australia (e-mails: i.pasquelmohottige@unsw.edu.au, h.habibi@unsw.edu.au, vijay@unsw.edu.au).

A. Vishwanath is with the IBM Research, Australia. (e-mail: arvishwa@au1.ibm.com).

S. S. Kanhere is with the School of Computer Science and Engineering, University of New South Wales, Australia. (e-mail: salil.kanhere@unsw.edu.au).

in post fire evacuation scenarios as a result of deactivating the HVAC system. We then quantify the power excursions of real evacuations across buildings of a university campus using occupancy data obtained from WiFi traces and a data-driven building thermal model. Our analysis demonstrates that evacuations (particularly during summer) can result in HVAC power excursions of up to 150% above the peak power demand threshold, imposing heavy power tariffs. Our **second** contribution develops a rule-based method for automatic rescheduling of "planned" evacuations retrospectively in order to eliminate the power excursions while adhering to evacuation standards of commercial buildings. **Finally**, we formulate an optimization problem for dynamic pre-cooling during reoccupation (especially during emergency evacuations) to minimize the collective cost of power excursion and delay in reaching to thermal comfort of occupants subject to maximum permissible power and temperature excursions (typically set by building managers). For a practical realization of our dynamic method, we relax the constraint (and associated cost) of occupants thermal comfort, and show (using real-world data from 43 evacuations) how power excursion is mitigated while incurring acceptable delays (about 5 minutes on average) in reaching to comfortable indoor temperature.

The rest of the paper is organized as follows: §II describes prior work on evacuation and energy management in commercial buildings, and §III describes our analysis that highlights the impact of evacuations on power consumption of building HVAC systems using real data. In §IV, we present a method for rescheduling planned evacuations to eliminate the power excursions, while in §V we develop a formal optimization for building pre-cooling post evacuation to minimize the cost of power and temperature excursions. The paper concludes in §VI.

II. Related Work

In this section, we first review related work on building evacuations and then discuss prior studies that employ occupancy derived by WiFi data for building energy management.

Building Evacuation: Related works on building evacuations are largely based on modeling and simulations but have not studied the role that WiFi session data of occupants can play in enhancing planning and execution of evacuations. Evacuations can either be planned (drills) or due to emergencies. Operational drills are routinely conducted to ascertain the efficacy of emergency evacuation procedures. In this context, work in [9] explored the possibility of applying network flow optimization models to building evacuations and work in [20] established an evacuation model that combined heuristic algorithms with network flow control to reduce evacuation time. Authors of [12] employed greedy algorithms to model building evacuations considering network flows that are constrained by the number of people present. In addition to drills, emergencies inside buildings, from fire to toxic chemical spillage, require immediate evacuation. Other scenarios, such as system faults or smoke from bush fires being sucked into the air ducts, may lead to false alarms, as reported in [25]. When such events occur, evacuations can last from a few minutes to a few hours. For example, a recent evacuation in Changi airport tower, Singapore [26] lasted for nearly 2 hours.

Occupancy Driven HVAC Energy Management in Buildings: While building energy management is a well-studied topic, an emerging body of work specifically uses WiFi-based occupancy information to reduce HVAC energy consumption. Melfi *et al*. [21] showed the potential of using existing IT infrastructure in buildings, including WiFi AP logs, to lower the building energy demand. Occupancy is implicitly obtained by tracking MAC and IP addresses at these APs, which in turn can be used to direct HVAC and lighting only to the occupied zones, thus saving energy. A practical system that infers occupancy using WiFi and uses it to control the HVAC of a commercial building leading to an 18% reduction in energy consumption, was demonstrated by Sentinel in [5]. Learning the spatial occupancy patterns enabled by WiFi connection logs and using that information to drive HVAC scheduling is shown to reduce the energy consumption of several buildings spanning a large campus by over 30% in work by Trivedi *et al*. [33]. The authors in [35] employed pre-cooling methods to reduce the HVAC energy consumption of commercial buildings. They employed "Gray-box" approaches to model the thermal dynamics of buildings, as opposed to physics-based "White-box" approaches [1], [4], [19] or purely data-driven "Black -box" approaches [32]. Inspired by work in [35], we use the Gray-box model, which accounts for occupancy and outside temperature. To the best of our knowledge, mitigating the adverse impact of HVAC power draw post emergency evacuations with the use of occupancy derived by WiFi session data has not been explored in the literature. Our work fills this important gap.

III. Impact of Evacuation Events on HVAC Power Consumption

In this section, we begin by introducing the buildings, data sets of the study, and the building thermal model. Next, our method to determine the peak demand threshold of a building is presented. Finally, we quantify the impact of evacuations on building HVAC power consumption.

A. Buildings and Datasets

This study was conducted on a large university campus, analyzing 43 evacuations that took place in 14 buildings. The buildings are primarily used for office and/or academic activity purposes. For our study we obtain two types of data sets; (1) building evacuation drill schedule and reports provided by campus estate management, and (2) daily WiFi session logs for all campus APs provided by the campus IT department. We note that appropriate clearance (UNSW Human Research Ethics Advisory Panel approval number HC190372) was obtained from the University ethics review board for this study. We further note that user IDs and MAC addresses contained in the WiFi logs are anonymized prior to storage/analysis by applying a one-way hash function.

Evacuation dataset: Campus Estate Management provided us with: (a) schedule of planned evacuations (aka drills) across all buildings on campus over the 6-month period (from Oct-2018 to May-2019) of this study, (b) copies of drill reports from seven planned evacuations over this period, and (c)

TABLE I
SAMPLE WIFI SESSION LOGS.

User ID	MAC addr	AP name	Assoc. time	Disassoc. time	Thput. (Kbps)
145e7e26	b76690ac	J17_F1_AP01	31/01/2019 10:40	31/01/2019 11:15	0.1
145e7e26	127d4fb7	J17_F1_AP02	31/01/2019 10:55	31/01/2019 11:20	8.6
b6c72a33	decf7837	J17_F3_AP12	31/01/2019 11:15	31/01/2019 11:20	561.8

date/time of all unplanned evacuations that were recorded for four of the buildings over the study period. Evacuation drills, typically one per year, are planned by Estate Management at the beginning of each calendar year. Each drill is attended by a team of fire wardens, some of whom walk through various floors to clear people out while others are stationed at the various exits. After the drill, a report is filed that estimates the timing aspects of the drill, such as time of alarm, end of evacuation, and start of reoccupation, as well as number of occupants evacuated. These reports are analyzed to check if adjustments need to be made to the evacuation procedures to speed it up, and to make qualitative judgments.

WiFi dataset: The campus IT department manages a rich WiFi infrastructure comprised of over 5000 access points (APs). We focus on 14 buildings and use the building occupancy profile derived from WiFi session data to develop a method to automatically detect evacuation events. The University IT department provided us with: (a) data showing the physical mapping of WiFi APs to buildings and floor levels, and (b) daily session logs across the 5249 APs for a period of 210 days from Oct-2018 to May-2019, from which we redacted data for the Christmas holiday period (since campus operations are minimal) to get 180 days of usable data. Combining the two gave us a total of around 65 million WiFi session records; Table I shows a representative snapshot. Each record contains a unique *User ID* (note that we have hashed this to preserve anonymity); device *MAC address* (also hashed to preserve anonymity); a unique *AP name* that clearly indicates the building name, floor level, and access point ID; time at which the device *associated to/disassociated from* the AP (note that this is in minutes and therefore we do not have sub-minute accuracy); and *avg throughput* indicating data rate during the session. For example, the top session in Table I belongs to user `145e7e26` who connected from device `b76690ac` to `AP1` located in building J17 floor F1, from 10:40am to 11:15am, and used an average throughput of 0.1 Kbps over the session. The second record represents the same user but with a different device `127d4fb7` connected to `AP2` located on the same floor of the building from 10:55am to 11:20am, while the third record is a different user connected to a different floor of the same building, connected for 5 minutes with an average throughput of 561.8 Kbps.

Temperature dataset: We obtained outdoor air temperature data (with hourly resolution) for the area of our university campus, spanning the period between Oct-2018 and May-2019, from a public repository of air quality and meteorological data provided by the New South Wales Department of Planning, Industry and Environment, Australia [27]. In Australia, with a moderate climate, the recommended set point for HVAC systems in summer is 25°C [6]. For typical commercial HVAC systems, the supply air temperature is set to its minimum value of 13°C [36]. We will apply our outdoor temperature dataset along with the desired set-point temperature and HVAC supply air temperature to compute the indoor zone temperature (§III-C) using the thermal models of individual buildings (§III-B).

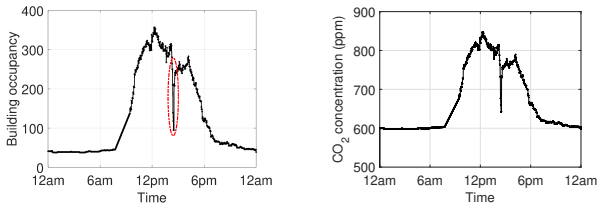

(a) Building occupancy profile obtained from WiFi data.

(b) CO_2 concentration computed from occupancy profile.

Fig. 1. Building daily profile of: (a) occupancy, and (b) CO2 concentration – note that a 20-min evacuation event occurs at 2:12pm.

TABLE II
NOMENCLATURE OF OUR GRAY-BOX MODEL.

Symbol	Description	Units
T_z	Zone temperature	C
T_s	Supply air temperature	C
T_a	Outside air temperature	C
$T_{z,sp}$	Building set-point temperature	C
C_z	Zone thermal capacitance	kJ/C
k_a	Heat transfer coefficient between ambient and zone	kW/C
$k_{o,2}$	Captures the non-occupancy loads such as plug loads and lighting	kW
θ_{CO2}	CO_2 concentration	$micromol/mol$
$k_{o,1}$	Maps the CO_2 concentration to internal heat gain	$kWmol/micromol$
$\dot{Q}_{cooling}$	Power demand of HVAC cooling	kW
$c_{p,a}$	Specific heat capacity of air	$kJ/kg/C$
$\dot{m}_s(t)$	Mass flow rate of conditioned air with temperature	kg/s
$\dot{m}_{s,o}$	Mass flow rate of air corresponding to the nominal position of the damper	kg/s
k_c	Effective gain term of the proportional controller	$kg/s/C$

B. Thermal Model of Building

We borrow the building thermal model from [35], and adapt it to individual buildings in our university campus. Specifically, we adopt the Gray-box approach to model evolution of zone temperature as a function of the heat gain from the outside, internal heat gain due to occupancy and zone cooling rate due to air renewal supplied by the HVAC. The mathematical representation of the Gray-box model is given by:

$$C_z \frac{dT_z}{dt} = \underbrace{k_a(T_a(t) - T_z(t))}_{\text{heat gain from outside air}} + \underbrace{k_{o,1}\theta_{CO2}(t) + k_{o,2}}_{\text{heat gain from occupants}} + \underbrace{\dot{Q}_{cooling}(t)}_{\text{cooling load}} \quad (1)$$

Here, C_z is the zone thermal capacitance which is without loss of generality assumed to be unity, as proposed in [38]. This means cooling load refers to the scaled cooling load and therefore, air conditioning power shown in the paper are scaled with respect to C_z. The heat gain from outside air is captured by the following terms: T_a, outside air temperature, T_z, zone temperature, and k_a, heat transfer coefficient between ambient and zone. The heat gain from occupants is modeled as an affine function where θ_{CO2} is CO_2 concentration, $k_{o,1}$ and $k_{o,2}$ are linear regression coefficients. $k_{o,1}$ maps the CO2 concentration to internal heat gain and $k_{o,2}$ captures the non-occupancy loads such as plug loads and lighting.

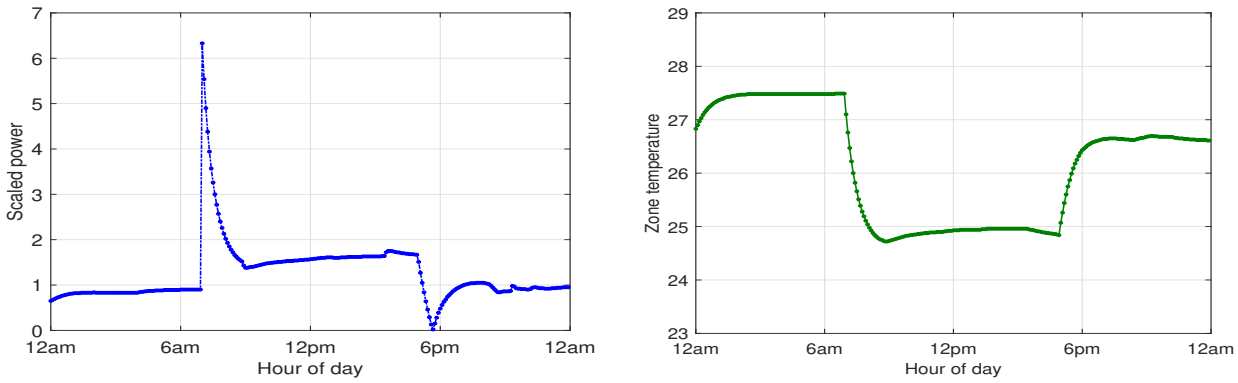

(a) Power demand. (b) Zone temprature.

Fig. 2. [Default Settings] daily trace of: (a) power demand, and (b) zone temperature; in building C22; on Tuesday 15-Jan-2019.

Fig. 1(a) shows the occupancy profile of building F25 during a day on which an evacuation occurred. The building occupancy is estimated by counting the total number of unique user ids connected to all WiFi APs in the building. To discount the impact of users who connect to the building APs while walking past, we filtered out the devices (from WiFi traces) with transient connections. Specifically, devices which record average daily session duration of less than 5 minutes in a building were not considered as occupants. The heat gain from occupants (CO_2 concentration) is obtained by computing the rate of change of occupancy. The corresponding CO_2 concentration is shown in Fig. 1(b). We capture the heat gain due to people exiting and reoccupying the building using the evacuation patterns observed in our previous work [23]. $Q_{cooling}$(t) in Eq. 1 is the rate of heat efflux attributed to cooling (cooling load) and it is given by Eq. 2.

$$\dot{Q}_{cooling}(t) = c_{p,a}[\dot{m}_s(t)(T_s(t) - T_z(t))] \tag{2}$$

$$\dot{m}_s(t) = \dot{m}_{s,o} + k_c(T_z(t) - T_{z,sp}(t)) \tag{3}$$

where $c_{p,a}$ is the specific heat capacity of air and $\dot{m}_s(t)$ is the mass flow rate of conditioned air with temperature T_s. As shown in Eq. 3, the mass flow rate is approximated to $\dot{m}_{s,o} + k_c(T_z(t) - T_{z,sp}(t))$ where $T_{z,sp}$ is the building set-point temperature. The coefficients $\dot{m}_{s,o}$ and k_c represent the mass flow rate of air corresponding to the nominal position of the damper (*i.e.,* a static offset in the actuator output) and effective gain term of the proportional controller, respectively. Table II summarizes nomenclature of our Gray-box model.

Deriving model parameters: It has been shown [17] that a data-driven thermal model (similar to our Gray-box model) developed for a building can be used to model the dynamics of indoor temperature in other buildings of similar size. In absence of Building Management System data for buildings of our university campus, we model the dynamics of indoor temperature using the scaled version of constant coefficients in Eq. 1 that were obtained from a commercial building studied in [35]. Note that, the accuracy of the approximated thermal model in our study cannot be evaluated without ground-truth data of indoor temperature. In what follows, we derive the scaled parameters ka, $k_{o,1}$, $k_{o,2}$, $\dot{m}_{s,o}$ and k_c based on physical properties (i.e., surface area, volume, and capacity) of individual buildings.

Considering Eq. 1, we note that in the first term $k_a = h.A$ (according to Newton's law of cooling), where h is heat transfer coefficient depending on materials of building (wall/roof/window), and A is the surface area through which the heat is transferred. We, therefore, scale this parameter k_a based on the buildings surface area with an assumption that heat transfer coefficients remain consistent across various buildings – buildings are made of the same materials. The parameters $k_{o,1}$ and $k_{o,2}$ in the second term of Eq. 1 correspond to heat gain of building occupants. We, therefore, scale them proportional to the occupant capacity of a building. The occupant capacity is computed by dividing the total floor area by unit area per person (*i.e.,* 5 m^2 in multi-purpose educational buildings [8]). Lastly, coefficients $\dot{m}_s$ and k_c in third term of Eq. 1 correspond to cooling load, and get scaled based on the total volume of the buildings. Note that the height of a floor in commercial buildings is assumed to be 3 m [31] when computing the surface area and volume of individual buildings.

C. Power Demand Threshold of Buildings

Having obtained the thermal model of individual buildings on our university campus, we now analyze the dynamics of cooling power and zone temperature during evacuation events. We note that the power excursion can only be quantified when the power demand threshold value of buildings is known. This threshold value is contracted between building managers and energy suppliers. In absence of such data in our study, we estimate the peak demand threshold value by analyzing the variation of daily power demand in our dataset.

For ease of illustration, let us consider an example building C22. Using the thermal model generated for this building (§III-B), we plot the time trace of power demand in Fig. 2(a) and zone temperature in Fig. 2(b) on a typical working day (Tuesday, 15 Jan 2019) in default settings without applying any control. The cooling load (power demand) is computed by Eq. 2 and the internal zone temperature (T_z) is recursively computed by Eq. 1. It is recommended [29] that HVAC systems are turned off during weekends and night hours of weekdays, and remain operational (tuned on) during days times (*e.g.,* 7am-7pm [2]) Monday to Friday. For our university campus, typical working hours are 8am-5pm, and hence HVAC

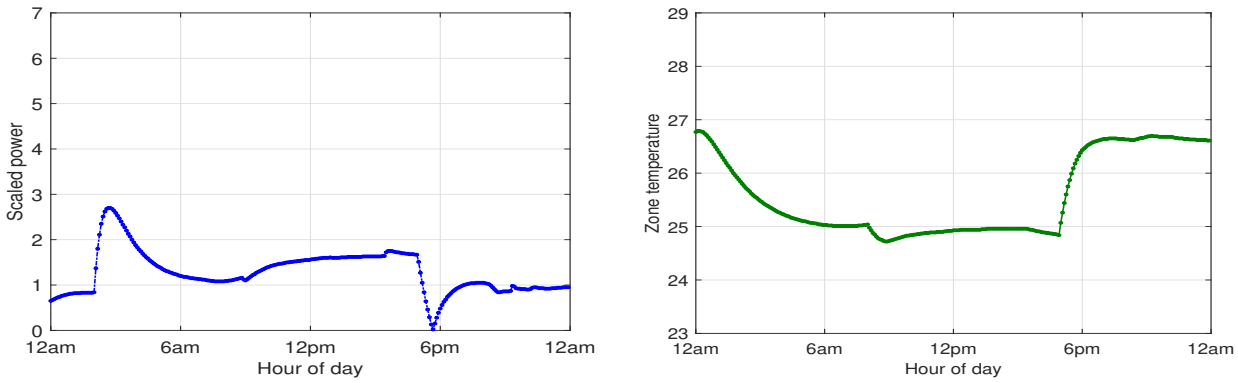

(a) Power demand. (b) Zone temprature.

Fig. 3. [Night Pre-cooling] daily trace of: (a) power demand, and (b) zone temperature; in building C22; on Tuesday 15-Jan-2019.

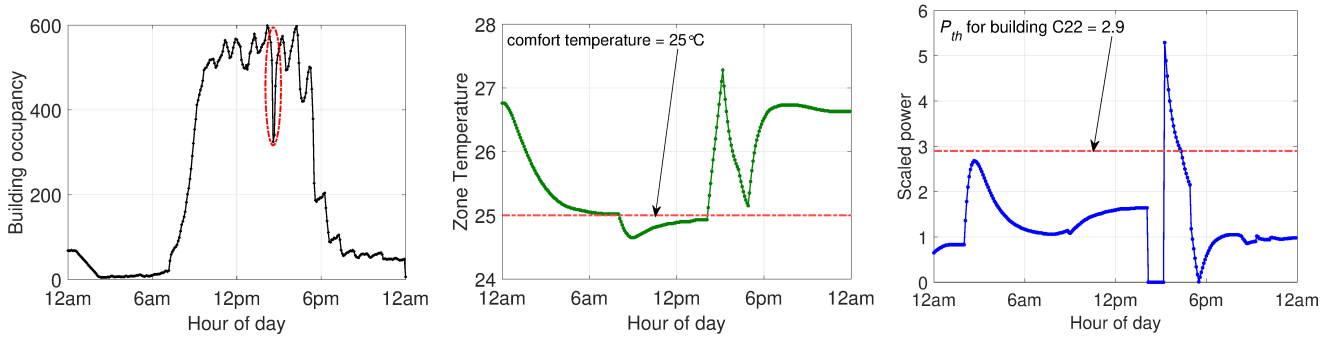

(a) Building occupancy. (b) Zone temperature. (c) Power demand.

Fig. 4. Daily trace of: (a) building occupancy, (b) zone temperature, and (c) power demand, for an evacuation drill that begins at 2:10pm and lasts for an hour; in building C22; on Friday 15-Mar-2019.

systems are turned on at 7am in the morning (an hour prior to expected occupancy time of 8am), so that comfort temperature is reached before occupants arrive, and turned off at 5pm while maintaining comfort conditions until the end of the expected occupancy. Therefore, in warmer months indoor zone temperature gradually increases (during evening/night hours) to levels that are several degrees higher than thermal comfort temperature – as shown in Fig. 2(b), indoor temperature has risen up to 27.5°C. In order to rapidly cool-down the building and ensuring a comfortable indoor temperature for occupants arriving at around 8am, the HVAC system starts full-load in the morning (between 7am and 9am), resulting in a significant peak of power demand, as shown in Fig. 2(a). As a result, indoor temperature steeply falls from 27.5°C to about 25°C, as shown in Fig. 2(b).

Note that, the power demand is a function of mass flow rate $\dot{m}_s(t)$, supply air temperature $T_s(t)$, and zone temperature $T_z(t)$, as expressed by Eq. 2. Ideally, when HVAC is turned off at 5pm, it is expected that the two impacting factors $\dot{m}_s(t)$ and $T_s(t)$ become zero, hence resulting in zero power demand. However, in our thermal model Eq. 1 (borrowed from [35]), $\dot{m}_s(t)$ approximates the rate of mass air flow (expressed by Eq. 3) and is a function of $T_z(t)$, and $T_(z, spn)(t)$, resulting in negligible values for power demand after 5pm before reaching to zero.

To avoid morning spikes in power demand, building managers often implement various pre-cooling strategies [29], [35] (circulation of cool air within a building) during the off-peak hours (night time/ early morning) with the intent of gradual cooling over a longer period prior to arrival of occupants.

It is important to note that we illustrate in Figures 2 and 3 the impact of two policies (default versus night pre-cooling) on the power consumption of a given building.

In this paper, we apply a heuristic-based night pre-cooling method (suggested in [35]) that starts after midnight every day depending on outside air temperature. Fig. 3 illustrates power demand and zone temperature for building C22 with pre-cooling applied on data of 15th of January 2019 in comparison to the Fig. 2 with default setting. The pre-cooling method employs Time Dependent Dijkstra (TDD) algorithm to bring down the zone temperature T_z to occupant thermal comfort temperature 25°C with minimal energy cost. We obtain the daily power demand profile of individual buildings with night pre-cooling applied across the 180 days. For each building, the maximum value of daily morning peaks is chosen as the power demand threshold of the building. We note that daily morning power demands are fairly consistent with a small standard-deviation of 0.13 (on average across all buildings). This suggests that for a given building, morning peak demand varies slightly during the course of 180 days. We denote the threshold peak power demand by P_{th} which will be used to quantify the power excursions in the following section.

D. Impact of Evacuation Event on HVAC Power Consumption

We now analyze 43 evacuation events (planned and emergency) that occurred in the 14 buildings and quantify their impact on the HVAC power consumption using thermal models

generated in §III-B. We quantify the magnitude of power excursions caused by an evacuation, as the maximum overshoot with respect to building P_{th} derived in §III-C. An evacuation event evolves in three stages: (i) immediate evacuation upon hearing the alarm, (ii) gathering at designated assembly point, and (iii) reoccupation. The HVAC is shut down during stages (i) and (ii), and we call this period as "evacuation duration" – during this period $Q_{cooling}$ is set to zero.

Upon commencement of reoccupation, the HVAC system is switched on, and hence the supply air temperature T_s is typically (default settings) set to its lowest set point 13°C in order to quickly bring the internal temperature to comfortable thermal levels.

Fig. 4 shows daily trace of occupancy, zone temperature, and power demand in building C22 for a day on which an hour-long evacuation drill took place – data points are shown with a 5-min resolution. The evacuation starts at 2:10pm as highlighted by a significant dip (from 600 to close to 300) in the building occupancy profile, shown in Fig. 4(a). The HVAC is automatically turned off at commencement of evacuation (2:10pm), and hence indoor temperature increase by 2.4°C (from 24.9°C to 27.3°C), as shown in Fig. 4(b). At reoccupation (3:15pm), the HVAC is turned on, attempting to rapidly cool down the building for occupants. As a result, the power demand exceeds P_{th}, resulting in a maximum overshoot of 82% shown in Fig. 4(c). Here, we quantify the magnitude of power excursions by the maximum overshoot (with respect to building P_{th}) caused by evacuation events, comparing various evacuations across buildings. In §V, we will accumulate power overshoots (when demand exceeds P_{th}) over time to quantify a total power cost incurred by evacuation events.

1) Power Excursion of Real Evacuations: We now look at the variation of HVAC power demand post-evacuation for all of the 43 evacuations in 14 buildings in our study and quantify their maximum overshoot.

Among the 43 evacuations, 33 of them resulted in power excursion (77%) – this is because our study spans months from Oct-2018 to Apr-2019 which corresponds to the warm/hot seasons in the Southern Hemisphere (Sydney, Australia). The magnitude of power excursions for these 33 evacuations ranges from 3% to about 140%. There are 3 evacuations with power excursions of more than 100%. The largest overshoot of 143% is for an evacuation in F21 on 06-Dec-2018 at 12:44pm followed by a 131% overshoot resulted from an evacuation in F25 on 31-Jan-2019 at 11:16am, and 114% overshoot resulted from an evacuation in F23 on 28-Feb-2019 at 1:43pm. It is interesting to note that all three of the above evacuations occurred during the middle of the summer months (Dec-Feb) with high daily outside temperature profiles. We show in Fig. 5 the building occupancy, outside air temperature and HVAC power demand for the evacuation with largest power overshoot in F21. We note that the evacuation duration is 105 minutes and the outside air temperature is 34.5°C. The power overshoot is less than 100% for the remaining 30 evacuations with power excursions. We note that for evacuations without power excursions the outside air temperature, T_a (ranging from $19.2°$ C - $23.1°$ C) and evacuation duration, ($\leq$30 minutes) were both lower than those values reported for evacuations with excursions.

TABLE III
IMPACT OF VARIOUS FACTORS ON POWER OVERSHOOT.

Factor	F-statistic	Mutual Information
Outside air temp.	0.63	1.00
Evacuation duration	1.00	0.75
Occupancy	0.12	0.21

2) Factors Impacting Power Excursion: Previously, we demonstrated that high power overshoot is reported for evacuations during summer months with high outside temperature values. In this section, we analyse how outside air temperature, evacuation duration and occupancy (at evacuation) correlate with the scale of power overshoot.

We compute uni-variate F-test statistic and mutual information to quantify the relationship between each of factors with power overshoot as shown in Table III. F-test captures linear dependency, while mutual information captures all types of dependencies. F-statistic rates evacuation duration as the most discriminating feature, closely followed by outside air temperature. Mutual information rates outside air temperature as the most discriminating feature, closely followed by evacuation duration. Both methods mark occupancy at evacuation as less related to power overshoot compared to the other factors. We plot in Fig. 6, the percentage power excursion against the outside temperature, evacuation duration and occupancy at evacuations. Fig. 6(a) and Fig. 6(b) show that higher the outside air temperature, and evacuation duration higher is the power overshoot. It seems that the quantitative measures largely agree with our intuitive perception.

IV. RESCHEDULING EVACUATION DRILLS

In the previous section, we demonstrated that excursions can occur in building HVAC power demand due to post evacuation cooling, incurring excessive energy costs. In this section, we aim to develop a method for rescheduling planned evacuations in buildings to manage their potential energy costs – note that the schedule of planned evacuations is known in advance. In order to comply with the requirements of building safety standards, evacuation drills are conducted at least once a year [24], [28]. We retrospectively analyze data traces of planned evacuations to see whether they can be rescheduled to eliminate the power excursion post evacuation, while fulfilling the safety standards, especially related to the number of occupants at the start of an evacuation.

A. Shifting Evacuation Events

Given the occupancy pattern of a planned evacuation, we retrospectively and synthetically shift the evacuation event to a different time and/or day, while relatively preserving the original occupancy pattern of the planned evacuation. In our previous work [23], we showed how evacuations display a sharp dip (*i.e.,* people evacuating) in the building occupancy followed by a sharp rise (*i.e.,* people re-occupying), forming a V-shape pattern. This V-shape pattern in building occupancy during a known planned evacuation is captured as the profile of the evacuation event ("evacuation pattern"). We reschedule an evacuation drill at a different time during the working hours of

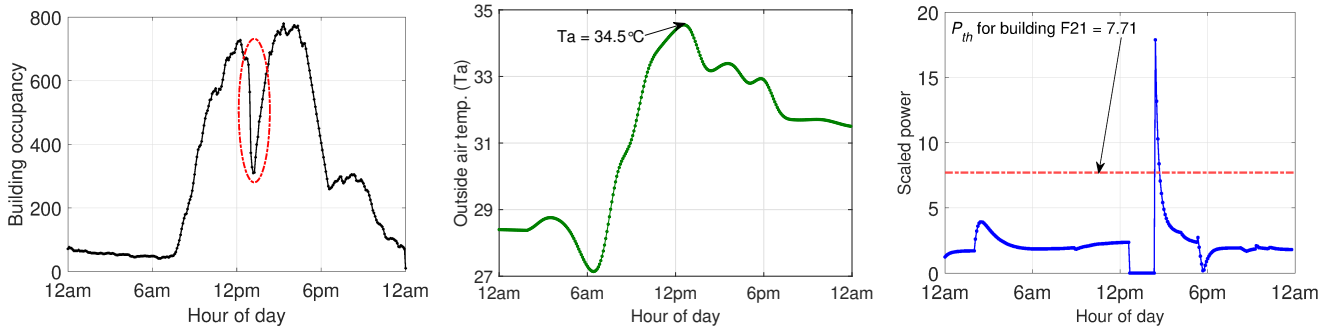

(a) Building occupancy. (b) Outside air temperature. (c) Power demand.

Fig. 5. Daily trace of: (a) building occupancy, (b) outside air temperature, and (c) power demand, for evacuation drill which begins at 12:44pm and lasts for an hour and 45 minutes in building F21

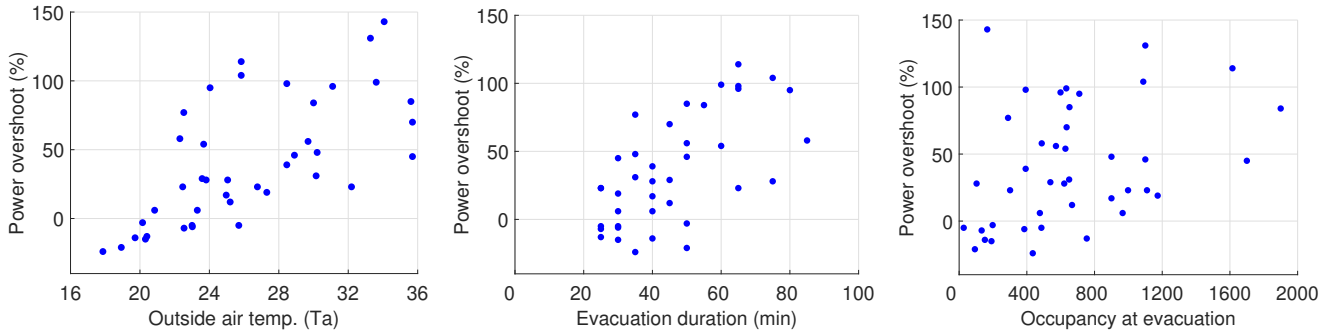

(a) Power versus outside air temperature. (b) Power versus evacuation duration. (c) Power versus occupancy at evacuation.

Fig. 6. Correlation of power overshoot with: (a) outside air temperature, (b) evacuation duration, and (c) occupancy at evacuation.

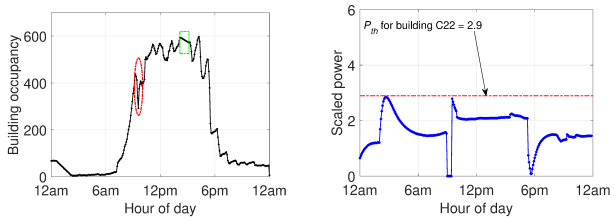

(a) Occupancy. (b) Power demand.

Fig. 7. Building daily profile of: (a) occupancy and (b) power demand, for an evacuation injected at 9am.

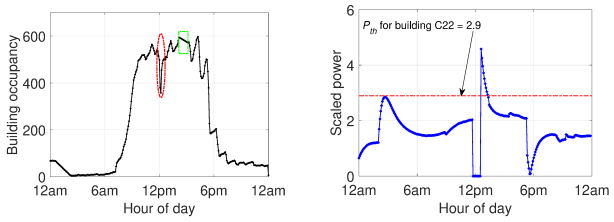

(a) Occupancy. (b) Power demand.

Fig. 8. Building daily profile of: (a) occupancy and (b) power demand, for an evacuation injected at 12pm.

days of the week in which the actual evacuation was originally planned, by temporal shifting (via simulation) its pattern across the building occupancy profile.

For the evacuation pattern, let there be n time epochs (of 5-minute) during an evacuation event, each occurs at time t_i with corresponding occupancy y_i where $1 \leq i \leq n$. Our objective is to inject a rescheduled evacuation at time $\hat{t}_1$ when the building occupancy equals to $\hat{y}_1$. This process requires us to adjust the building occupancy for n epochs (starting from $\hat{t}_1$) according to dynamics of occupancy during the original planned evacuation. Eq. 4 is used to adjust the building occupancy (*i.e.,* $\hat{y}_2$ to $\hat{y}_n$) for the shifted evacuation.

$$\hat{y}_{i+1} = \hat{y}_i \left(1 + \frac{y_{i+1} - y_i}{y_i}\right) \tag{4}$$

Once the evacuation is shifted to a new time $\hat{t}_1$, it is necessary to fill the gap in the occupancy profile corresponding to the duration of actual planned evacuation. We construct new occupancy data points by applying linear interpolation connecting the first and last data-points of the gap in the occupancy profile. Following this update to the daily occupancy profile (due to rescheduled evacuation), the power demand is recomputed using the thermal model (Eq. 1) and modified occupancy.

To illustrate the process, let us start from Fig. 4(a) where the occupancy profile of building C22 is shown with an evacuation scheduled at 2:10pm. We show in Figures 7 and 8 how this evacuation is synthetically shifted to 9am and 12pm, respectively on the same day. It is seen that the evacuation rescheduled at 9am does not cause a power excursion (daily power never exceeds the building threshold power $P_{th} = 2.9$ in Fig. 7), while the evacuation rescheduled at 12pm results in a 58% excursion at 12:40pm as shown in Fig. 8. Note that the original evacuation pattern is patched by linear interpolation (highlighted by green box) in both Figures 7 and 8.

B. Typicality of Building Occupancy

An important parameter pertinent to an evacuation drill is the number of building occupants at the scheduled start time of

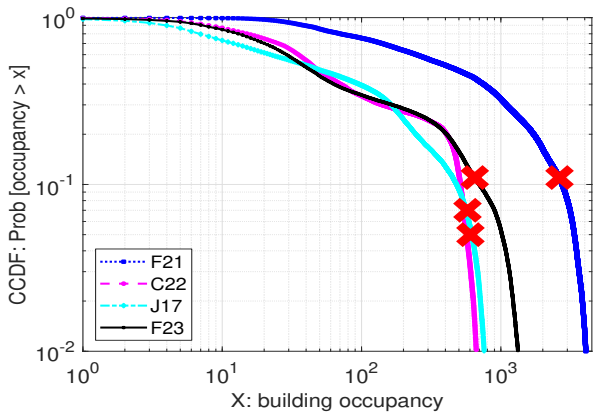

Fig. 9. CCDF of building occupancy (number of connected devices) during 6 months, highlighting (red × markers) the occupancy at planned evacuation.

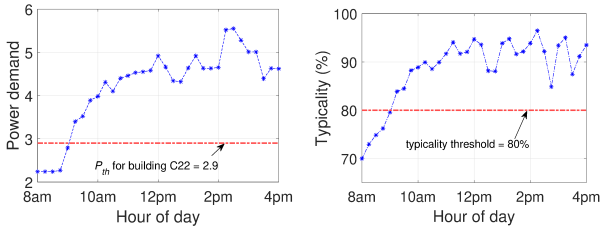

(a) Maximum power demand. (b) Typicality.

Fig. 10. Variation of (a) maximum power demand, and (b) typicality, for evacuations rescheduled on the same day of the original planned evacuation in building C22.

drill. Regulations for building fire safety [7] recommend that fire drills should be conducted with "appropriate number" of people inside a building. This is indeed merely a qualitative measure. We, instead, quantify it by using the "typicality" metric – this is the percentile level at which the building is occupied at start of evacuation compared to its historical occupancy count [23]. This quantification helps to specify measurable guidelines for scheduling drills. Fig. 9 shows the CCDF of building occupancy in four representative buildings from our dataset. We mark (red ×) the occupancy of each building at the start of the planned evacuation. It can be seen that the typicality is fairly high in percentile (seemingly appropriate) across the representative buildings with F21 at 89%, C22 at 93%, J17 at 95%, and F23 at 89%. Note that all 14 planned evacuations demonstrated high typicality highlighting that building managers intentionally schedule them at times when buildings are relatively highly occupied ($> 80\%$), adhering to standards.

C. Rescheduling Evacuation Drills

We now develop a method to reschedule the planned evacuation in order to eliminate the costs of power excursions, while ensuring the typicality measure of building occupancy is also appropriate.

We synthetically shift (via simulating data traces) a rescheduled evacuation event during working hours (*i.e.,* 8am-4pm). At each run of the simulation, we shift the start of the evacuation ahead by 15 minutes, recording the maximum power demand due to cooling post evacuation as well as the typicality of building occupancy. The objective is to ensure

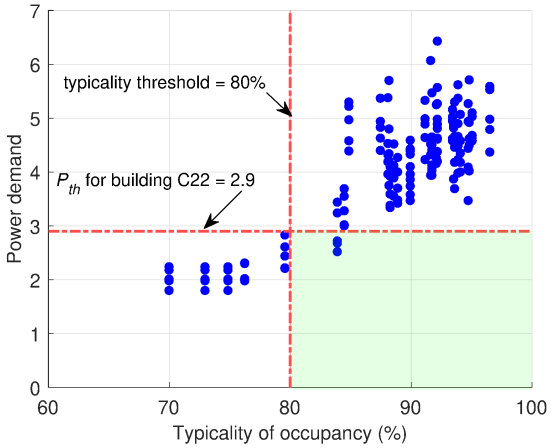

Fig. 11. Maximum power demand versus typicality of rescheduled evacuations (injected every 15 minutes) during the week of the planned evacuation.

Algorithm 1 Finding desirable rescheduled evacuation drills.

$A_T \leftarrow \{t_1, t_2, t_3, ..., t_N\}$
  where t_i's are chosen at 15-min resolution during working hours in the week of planned evacuation
$P_{th} \leftarrow P^*$
$T_{th} \leftarrow 80\%$

function RESCHEDULEDTIME (A_T, P_{th}, O_{th})
  $A_{ST} \leftarrow \{\}$
  for t belong to A_T **do**
    $P_{max} \leftarrow$ *compute max power demand at* t
    $T_m \leftarrow$ *compute typicality measure at* t
    if $P_{max} < P_{th}$ and $T_m > T_{th}$: **then**
      $A_{ST}.insert(t)$
  return A_{ST}

that: (a) the maximum power demand remains below P_{th} of the building (contracted with the energy provider), and (b) typicality is more than a threshold, say, 80%.

We aim to find those rescheduled evacuation events that satisfy the above objectives. In Fig. 10, we plot dynamics of maximum peak power demand and typicality for evacuations rescheduled every 15-minute for the building C22 on the same day as the original planned evacuation. The dotted red line in Fig. 10(a) highlights the building peak power threshold P_{th} that are rescheduled between 8am and 4pm. We observe that the maximum power demand due to evacuation is higher than the building P_{th} for the simulated evacuations. As observed from Fig. 10(b), the typicality during evacuation consistently remains above the threshold of 80% for all simulated evacuations between 9am to 4pm. Therefore, none of the rescheduled times (on the same day as the original planned evacuation) meet both of the requirements, namely maximum power demand and typicality of occupancy. Note that the search algorithm is repeated for remaining days of the week during which the original evaluation was planned, to minimize disruptions to scheduled planning, although we could have technically considered rescheduling to the week before and after. The pseudo code of the method proposed to determine potential times to reschedule the planned evacuation

is shown in Algorithm 1.

Fig. 11 shows the scatter plot of maximum of power demand versus typicality for all potential rescheduled evacuation times (at 15-min resolution) during the entire week (from 11-Mar-2019 to 15-Mar-2019) during which the drill in C22 was originally planned. Note that each blue dot on this plot represents a potential rescheduled evacuation. We show power demand threshold with horizontal dotted red lines; it is the maximum of the daily morning peaks of the building across 180 days (with pre cooling), and the minimum requirement for typicality of occupancy with vertical red dotted line. The bottom right quadrant shaded in green, represents the instances when both requirements are satisfied and is thus the region of interest. We observe that there are three different time epochs which match our desired criteria.

We, in §III, analyzed 14 planned evacuations and showed that 10 of them incurred post evacuation power demand excursions. For those 10 buildings, we run our search algorithm to find a desirable time-slot for rescheduling the evacuation drill. Our algorithm finds an appropriate time to reschedule 7 evacuations within the same week. However, for evacuation drills in three buildings, namely F21, C20, and F25, it was impossible to reschedule them at a different time in the week they were originally planned. We found for buildings C20 and F25 that the average of their daily occupancy profiles during our search within the week was respectively, 25% and 38% lower than their average occupancy across the entire 6-month period, resulting in low typicality, which explains the inability of our search algorithm to find a suitable epoch for rescheduling the evacuation. Note that the typicality of occupancy at actual planned evacuations as it took place for buildings C20 and F25 are closer to the typicality threshold with values 86% and 82% respectively. For building F21, the average temperature during the week under consideration was 5.6°C (25%) warmer than the average temperature over the entire 6-month period, consistently causing power excursions, which explains why rescheduling is not feasible. We also found that the temperature was relatively high during the week before (∼10% warmer than average) and after (∼20% warmer than average) the originally planned week. One can extend the search horizon beyond one week (before and/or after), but this is outside the scope of this paper.

V. Optimizing Cost of Energy and Thermal-Discomfort Post Evacuation

In the previous section, we demonstrated how evacuation drills can be (retrospectively) rescheduled for another time and/or day during the week they are originally planned to eliminate the undesirable energy costs due to HVAC power demand excursions. We also showed for some buildings it may not be always possible to find a suitable time slot to reschedule drills. However, it is important to note that rescheduling cannot be considered as a solution for unplanned evacuations. Therefore, in this section our objective is to find a way to minimize the post-evacuation cooling energy costs while the thermal comfort of occupants is maximized.

A. Typical Behavior of HVAC System Post Evacuation

When HVAC is activated post evacuation, it typically runs at full capacity to rapidly cool down the building. We plot in Fig. 12(a) the dynamics of power demand when a static cooling T_s is applied for a representative planned evacuation which occurred during morning time on 10-Apr-2019 in building M15. For comparison, we show how the power demand changes when the static set point varies from its lowest possible value ($T_s = 13°$C shown by dotted blue line) to slightly larger values (15°C and 17°C, respectively shown by pink line with cross markers and black line with star markers).

When T_s is set to 13°C, the building HVAC system runs with its full capacity to rapidly cool down the building, resulting a maximum excursion of 38%. Increasing the set point T_s to 15°C would halve the maximum excursion, lowering it to 17.3%. A further increase in T_s, setting it at 17°C would cause the power demand to fall under P_{th}, and thus mitigating the excursion. Fig. 12(b) illustrates the building zone temperatures corresponding to these three set points. We observe that the comfort temperature 25°C is achieved within an hour from the commencement of building re-occupation. When Ts = 15°C, thermal comfort is reached in 90 minutes. However, when Ts = 17°C, the comfort temperature is only reached just prior to close of business at 4pm. The key message from this observation is that the power demand can be appreciably controlled by lowering the cooling load (reducing the HVAC set point T_s), but at a cost of delaying the time taken in reaching the comfort temperature with the caveat that excessive delays are undesirable for building occupants.

B. Optimizing Energy and Thermal Comfort

We now develop a non-linear optimization problem for post evacuation cooling to minimize the cost associated with (a) building power excursion, and (b) occupants thermal comfort, subject to certain constraints. The optimization framework allows us recompute the optimal value of HVAC supply air temperature dynamically (every epoch time, say 5-minute). Our objective is to minimize the total cost at each epoch given by:

$$\min\ J(k) = \alpha.P_E(k) + \beta.T_E(k) \tag{5}$$

where $\alpha.P_E(k)$ is power excursion cost and $\beta.T_E(k)$ is occupants discomfort cost at each epoch time k, where $k = 1, 2, \ldots, N$. Note that N denotes the number of time epochs between start of reoccupation and end of working hours. Scaling parameters α and β are constant. The cost $J(k)$ is optimized independently at every epoch. $P_E(k)$ is the power overshoot with respect to the building demand threshold P_{th}, and is computed by:

$$P_E(k) = \dot{Q}_{cooling}(k) - P_{th} \tag{6}$$

where $\dot{Q}_{cooling}$ is the power demand of HVAC cooling given by:

$$\begin{aligned}\dot{Q}_{cooling}(k) = c_{p,a}(\dot{m}_{s,o} + k_c(T_z(k) - \\ T_{z,sp}(k)))(T_s(k) - T_z(k))\end{aligned} \tag{7}$$

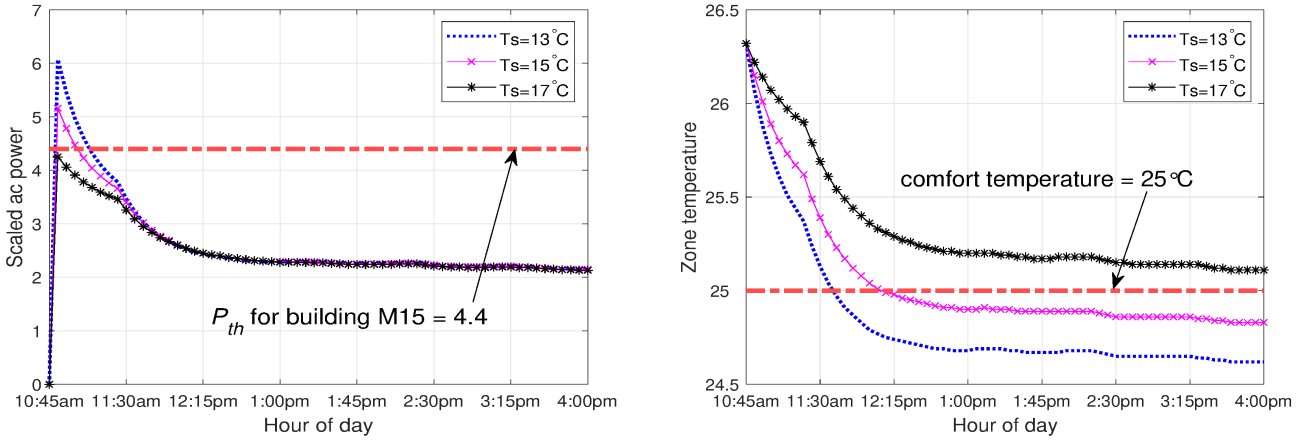

(a) Power demand.

(b) Zone temperature.

Fig. 12. Time trace of: (a) power demand, and (b) zone temperature, post-evacuation given static values of supply air temperature (T_s) in building M15.

To compute thermal discomfort, there exist some standard metrics such as Predicted Mean Vote (PMV) and Predicted Percentage Discomfort (PPD) [11], [37]. PMV is an index that estimates the mean of the thermal sensation vote from a group of people while PPD estimates the percentage of thermally dissatisfied people from PMV. For our study, we did not have access to a range of data such as mean radiant temperature, relative air velocity, relative humidity, metabolic rate and clothing insulation, which are required for PMV/PPD estimation, and hence considering these thermal comfort metrics is beyond the scope of our paper. Instead, we compute $T_E(k)$ as the overshoot of comfort temperature with respect to $T_{min}(k)$, the lowest achievable zone temperature when $T_s = 13°C$, and is computed by:

$$T_E(k) = T_z(k) - T_{min}(k) \tag{8}$$

where zone temperature, T_z, is a function of supply air temperature T_s, and is computed by:

$$C_z \frac{T_z(k+1) - T_z(k)}{\Delta T} = k_a(T_a(k) - T_z(k)) + k_{o,1}\theta_{CO2}(k) + k_{o,2} + \dot{Q}_{cooling}(k) \tag{9}$$

where ΔT is the control interval (say, 5 minutes) for setting the supply air temperature. Our decision variable is the supply air temperature of the HVAC system, T_s, in order to obtain an optimal cooling while the total cost of energy and thermal comfort is minimized. We note that T_s is capped by the lowest supply air temperature, ${T_s}^*$ (typically 13°C) that provides the maximum cooling, and hence our first constraint is given by:

$$T_s(k) \geq {T_s}^* \tag{10}$$

Further, we assume that relative excursions of power and thermal comfort are respectively capped at ${P_E}^*$ (say, 10%) and ${T_E}^*$ (say, 1%), as per requirements of the building manager, giving two more constraints:

$$\frac{P_E(k)}{P_{th}} \leq {P_E}^* \tag{11}$$

TABLE IV
NOMENCLATURE OF OUR OPTIMIZATION FRAMEWORK.

Symbol	Description (at time instance k)	Units
$P_E(k)$	Power overshoot with respect to P_{th}	kW
P_{th}	Building power demand threshold	kW
${P_E}^*$	Maximum power overshoot (configured)	–
$T_E(k)$	Overshoot of comfort temperature with respect to $T_{min}(k)$	C
T_{min}	Lowest achievable zone temperature capped by T_s^*	C
${T_E}^*$	Maximum overshoot in comfort temperature (configured)	–
$T_s(k)$	Supply air temperature	C
${T_s}^*$	Lowest supply air temperature	C
β	Scaling parameter associated with power excursion	$\$/C$
α	Scaling parameter associated with thermal comfort	$\$/kW$

$$\frac{T_E(k)}{T_{min}} \leq {T_E}^* \tag{12}$$

Table IV summarizes nomenclature of the optimization framework.

Computing Constant Parameters α and β: We now describe how scaling parameters α and β are determined to compute the cost of power excursion and occupants discomfort. The coefficient α is the cost incurred per unit of power excursion. We discussed in §I that peak demand charge is an essential component of the monthly electricity bill for commercial/industrial subscribers who typically generate a higher peak load than residential subscribers. Authors of [35] reported that a static penalty of ≈A$25 is applied per unit power excursion for a medium-size commercial building when the maximum demand recorded over a billing period exceeds the peak demand threshold. We compute α for our buildings by scaling the cost per unit excursion given by [35], proportional to the volume of buildings, similar to the way we obtained the thermal model of individual buildings in §III-B. There are other strategies that dynamically adjust the unit cost per excursion in proportional to the amount of overshoot beyond the peak demand threshold [13].

Moving to the scaling parameter β, the cost incurred per unit of temperature excursion, we note that commercial buildings invest some money on an annual basis to improve the quality of work environment. Indoor thermal comfort is one of the key factors determining the quality of work environment. We,

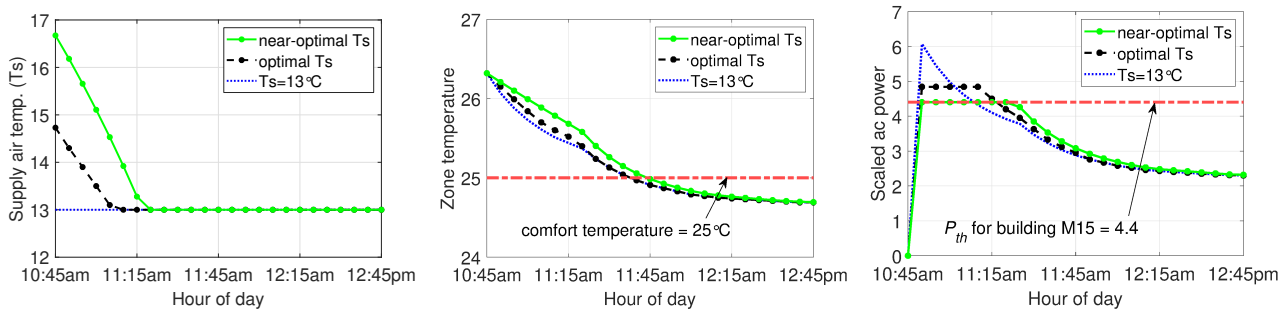

(a) Supply air temperature. (b) Zone temperature. (c) Power demand.

Fig. 13. Time-trace of dynamics of: (a) power demand, (b) temperature, and (c) supply air temperature, post-evacuation when T_s is set dynamically, for an evacuation in building M15 that lasts 35 minutes.

therefore, use this annual investment as a proxy of discomfort cost if the indoor temperature goes above the standard comfort level. Let us assume that organizations spend about 25% of their personnel costs in improving the quality of work environment. According to the annual report of our university in 2019 [34], the average salary cost of each employee is nearly A$180,000 per annum. Given a total of 6700 full-time employees, the total annual "employee cost" becomes A$1300 million. Therefore, considering 40 working hours per week and 52 weeks per year, the investment of thermal comfort per employee per each 5-minute epoch becomes equal to A$2. The annual report further states that the ratio of students to staff is 9 to 1 in our university. We assume that the capacity of each building is shared by students and staff uniformly with the same ratio, and hence compute the parameter β for each building as A$2 $\times N_{staff}$, where N_{staff} is the staff fraction (*i.e.,* 10% in our scenario) of total building capacity.

It is important to note that α and β are constant scaling parameters provided as input to our optimization. The value of these parameters is determined based on certain contextual values like dollar penalty per unit of power excursion and personnel cost. The above values are illustrative figures of α and β for our specific context that are obtained from public data of our country and university. One may choose a different method and/or data for setting these input values (adjusting the weight of energy cost against thermal comfort) before applying the optimization.

C. Practical Optimal Cooling Strategy

We formulated a non-linear optimization problem in the previous subsection to achieve a desired level of thermal comfort for occupants with reduced energy costs. Building facility managers may not have the resources and know-how to practically execute the non-linear optimization described above. To enable ease of use of the optimal framework, we propose a heuristic algorithm to find a "near-optimal" solution for this problem. To do so, we relax the thermal comfort condition and aim at eliminating the energy costs of power excursion due to evacuation. Obviously, T_s will be adjusted dynamically during 5-min epochs (relatively short intervals), so that the peak demand never exceeds P_{th}.

Our method for a near-optimal solution is as follows. We search for the best T_s value every epoch by initializing $T_s = {T_s}^*$, where ${T_s}^*$ is the lowest supply air temperature (typically 13°C) providing the maximum cooling. If the resulted power demand (computed by Eq. 2) exceeds the threshold P_{th}, yielding an overshoot of γ percent, then we increase T_s by the same factor γ. Note that the rate of change of temperature is linearly correlated with change rate of HVAC cooling load according to Eq. 2, and hence one may use a fraction of γ to adjust T_s. This search process continues until the resulting power falls below the threshold P_{th}. We note that our strategy may become too conservative, resulting in a significant power undershoot ($\gamma < 0$) – in that case, we need to reduce T_s accordingly again proportional to γ. Our search terminates when $\mid \gamma \mid < 0.01$. The pseudo code of the algorithm with asymptotic time complexity of $O(n)$ is shown in Algorithm 2.

Algorithm 2 Computing supply air temperature dynamically.

```
1:  T_s ← T_s*
        where T_s* is lowest supply air temperature.
2:  P_th ← P*
        where P* is threshold peak power demand.
3:  function ComputeSupplyAirTemp (T_s)
4:      P(T_s) ← Q̇_cooling at T_s
5:      γ ← (P(T_s) − P_th)/P_th, γ is power overshoot.
6:      if γ > 0: then
7:          T_s ← (1 + γ) * T_s
8:          return ComputeSupplyAirTemp (T_s)
9:      else if |γ| < 0.01 then:return T_s
10:     else
11:         T_s ← (1 − γ) * T_s.
12:         return ComputeSupplyAirTemp (T_s)
```

D. Evaluation of Optimization Methods

In this section we evaluate the non-linear optimization (optimal solution) and the proposed practical cooling strategy (near-optimal solution) on a planned evacuation (discussed earlier in Fig. 12) and compare the two methods.

We show in Fig. 13 how supply air temperature (T_s) , zone temperature (T_z), and power demand vary post evacuation for the planned evacuation in building M15. The evacuation causes a power excursion of 38% as shown in blue dotted line in Fig. 13(c). In §III we showed that this can be rescheduled within the same week as originally planned, such that there is no excursion in power. Fig. 13(a) illustrates the dynamics of T_s adjustment obtained from our near-optimal solution (shown

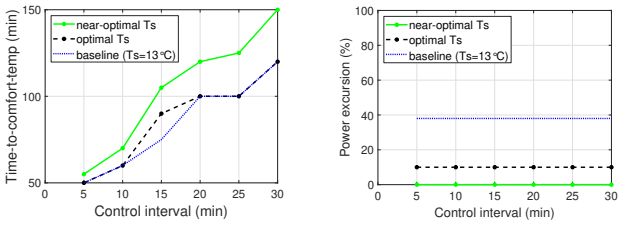

(a) It takes longer to reach thermal comfort as control interval increases. (b) Power excursion is independent of the control interval.

Fig. 14. Impact of control interval of T_s on thermal comfort and peak power demand.

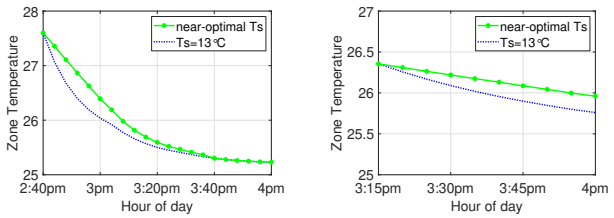

(a) Bldg F21, evacuation at 2:01pm. (b) Bldg C20, evacuation at 2:52pm.

Fig. 15. Samples of post evacuation where the zone temperature cannot reach to 25°C by 4pm (end of working hours) with our optimal cooling.

TABLE V
SUMMARY OF METHODS (THEIR USE, OUTCOME, AND COMPLEXITY).

Proposed solution	Applicable Domains	Power excursion	Thermal discomfort	Time shift	Time complexity
Rescheduling (Rule-based heuristic)	Drills only	Zero	-	0-120 hours	Low (<1 sec per run)
Optimal (Non-linear optimization)	Drills & Emerg.	Minimized	Minimized	Zero	High (≈ 3 min per run)
Near-optimal (Practical optimization)	Drills & Emerg.	Zero	Non-Zero	Zero	Low (<1 sec per run)

by solid green line) compared to the optimal solution obtained by brute forced method (shown by dashed black line) and the baseline scenario where T_s is statically set to its lowest value ${T_s}^* = 13°C$ (shown by dotted blue line). It is seen that our near-optimal solution cannot start the supply air temperature below 16.7°C in order to avoid the power excursion, but the optimal solution starts at 14.7°C which minimizes the total cost. It takes 35 minutes for our near-optimal T_s to reach to its maximum cooling value ${T_s}^*$, while this measure is 25 minutes for the optimal solution.

As a result, with our near-optimal strategy, a comfortable indoor temperature (25°C) is achieved with a fairly low delay (5-minute) compared to the baseline scenario, as shown in Fig. 13(b). Such delay is not observed with the optimal solution since up to 10% power excursion is allowed for the optimal solution – note that our near-optimal solution aims to mitigate the power excursion. We note that brute force search needs on average more than 3 minutes per each run (epoch) to compute the optimal solution in a search space of a billion data points. The search space would be much larger if upper bounds of ${P_E}^*$ and ${T_E}^*$ are chosen more generously, and hence a longer time is needed to find the optimal solution.

On the other hand, our near-optimal algorithm takes less than 1 second to produce the near-optimal solution per epoch, which makes it practical for dynamic cooling in real-time. Note that our near-optimal algorithm was implemented in Python running on an Intel i7 machine with 16 GB memory. During each time epoch, the algorithm converges after performing tens of iterations within a second. The algorithm searches for the best (lowest) T_s that results in power demand of maximum overshoot $\pm 0.01\%$. Importantly, we observe in Fig. 13(c) that our near-optimal setting of T_s completely mitigates the power excursion, while both the optimal and baseline solutions incur power excursions of 10% and 38%, respectively.

Evaluation of Dynamic Cooling for Real Evacuations: We evaluated the efficacy of our near-optimal dynamic cooling strategy for the 33 evacuations which resulted in some level of power excursion. We compute the delay (compared to the baseline) in reaching the comfort temperature during the working hours (8am-4pm). Note that for seven evacuations that occurred in the afternoon the indoor temperature does not reach to the comfort level 25°C prior to the end of working hours even with the baseline cooling strategy. Fig. 15 illustrates two examples as follow: (a) for four cases, our dynamic method manages to bring the indoor temperature to the level of the baseline method by 4pm, as shown in Fig. 15(a), hence is considered as zero delay, and (b) for three cases, our method can only meet the baseline temperature post 4pm, as shown in Fig. 15(b), hence are excluded. The average delay across, the 30 evacuations (in which our method meets the comfort temperature or baseline level), was found to be less than 7 minutes which is acceptable. The largest delay of 15 minutes was recorded for the largest building F21. Our proposed practical strategy can be incorporated into existing HVAC systems. It can be configured to activate post evacuation automatically. Note that in our evaluation, we relaxed the constrain of thermal comfort to mitigate power excursion costs. That said, to balance the thermal comfort of occupants against the associated cost of power excursions, one needs to solve the formal optimization problem as formulated in Eq.5 with all constrains included – mathematical analysis of trade-off is beyond the scope of this paper.

Impact of control interval: We have so far used a fixed control interval ($\Delta T = 5$ minutes) which is the resolution of our occupancy data. We now experiment with larger values of the control interval and quantify how they affect the time to reach to the thermal comfort temperature. Note that, the resolution of our occupancy data does not allow us to experiment with values less than 5 minutes. In Fig. 14, we show how the amount of time taken to reach thermal comfort temperature (time-to-comfort-temp), and power excursion varies by different control intervals, considering our three solutions, namely baseline, near-optimal, and optimal, for a representative building in our study. As shown in Fig. 14(a), the time-to-comfort-temp increases with control interval for all the three solutions. This is because the supply air temperature is adjusted very infrequently (a conservative approach of larger control intervals), which causes the zone temperature to fall slower compared to frequently adjusting the supply air temperature (an aggressive approach with smaller control intervals). It can also be seen that the gap between near-optimal and baseline solutions is widened for greater values of the control interval. As observed in Fig. 14(b), power excursion is unsurprisingly independent of the control interval. This is because the optimal solution aims to cap the power excursion at 10%, while our near-optimal solution is designed to mitigate the power excursion. Finally, the baseline approach incorporates no control.

Summary of Methods for Reducing HVAC Power Consumption: In Table V, we present a comparison summary of the three methods developed in this paper. Our rescheduling method is only applicable to planned evacuations where building managers have control over their scheduled time. By retrospectively shifting an evacuation to an appropriate time within the same week in which it was originally planned, our method aims to eliminate the power excursion associated with post evacuation HVAC cooling. On the other hand, our non-linear optimal and practical optimal methods for indoor cooling are applicable and adaptable to both types of evacuations, planned and emergency. Regardless of the time of occurrence of evacuations, our optimization method dynamically adjusts the supply air temperature using the thermal model of the building. Our proposed near-optimal cooling strategy (quantified across 33 emergency evacuations) reduces the computational complexity (at the expense of slight thermal discomfort) in managing power demand post evacuation, and hence is more practical to be used in real-time.

VI. Conclusion

In this paper, we first showed that evacuation events give rise to building power excursions and then quantified the power overshoot with respect to the peak power demand. Secondly, we developed a method for rescheduling evacuation drills to eliminate the power excursions that occur as a result of cooling post evacuation. Thirdly, we developed an optimization framework to minimize total costs associated with both power excursion and occupant thermal discomfort subject to supply air and maximum overshoot constraints. Lastly, we proposed a near-optimal cooling strategy that constrains the power consumption by dynamically adjusting the HVAC supply air temperature during reoccupation. In future work, we plan to consider the trade-offs between loss of productivity and cost of cooling in scenarios where occupants may have to return to the building the following day due to smoke accumulated indoors.

References

[1] F. Amara *et al.*, "Comparison and Simulation of Building Thermal Models for Effective Energy Management," *Smart Grid and Renewable Energy*, vol. 6, no. 7, pp. 95–112, Jan 2015.

[2] Australian Fair Work Obbudsman, "Hours of work," https://www.fairwork.gov.au/employee-entitlements/hours-of-work-breaks-and-rosters/hours-of-work, 2020, [Online; accessed 21-07-2020].

[3] Australian Government, "Heating, Ventilation and Air Conditioning," http://eex.gov.au/technologies/heating-ventilation-and-air-conditioning/ ., [Online; accessed 21-01-2020].

[4] P. Bacher and H. Madsen., "Identifying Suitable Models for the Heat Dynamics of Buildings," *Energy Build.*, vol. 43, no. 7, pp. 1511–1522, Jul 2011.

[5] B. Balaji *et al.*, "Sentinel: Occupancy based HVAC actuation using existing WiFi infrastructure within commercial buildings," in *Proc. ACM SenSys*, Roma, Italy, Nov 2013.

[6] Brendon O'Neill, "What is your air con temperature setting costing you?" https://www.canstarblue.com.au/electricity/air-con-temperature-costing/, 2019, [Online; accessed 13-01-2020].

[7] Canada Government, "Fire Drills," Office of the Fire Marshal and Emergency Management., Fire Drill Guideline, 2016.

[8] Certified Commercial Property Inspectors Association, "Commercial Property Safety Requirements: Maximum Occupancy," https://ccpia.org/occupancy-load-signs/, 2020, [Online; accessed 20-02-2020].

[9] L. Chalmet *et al.*, "Network models for building evacuation," *Fire Technology*, vol. 18, no. 1, pp. 90–113, Feb 1982.

[10] D. Chen *et al.*, "What the Indoor Air Temperatures in Houses in Three Australian Cities Tell Us," in *Proc. Windsor Rethinking Comfort*, Windsor, UK, 4 2018.

[11] T. Cheung *et al.*, "Analysis of the accuracy on PMV – PPD model using the ASHRAE Global Thermal Comfort Database II," *Building and Environment*, vol. 153, pp. 205 – 217, Apr 2019.

[12] W. Choi *et al.*, "Modeling of building evacuation problems by network flows with side constraints," *European Journal of Operational Research*, vol. 35, no. 1, pp. 98–110, Apr 1988.

[13] CIRCUTOR, "How to avoid maximum demand penalties in your electricity bill," http://circutor.com/en/documentation/articles/973-how-to-avoid-maximum-demand-penalties-in-your-electricity-bill/, 2015, [Online; accessed 24-06-2020].

[14] Clean Energy Group, "An Introduction to Demand Charges," National Renewable Energy Laboratory, Fact Sheet, Aug 2017.

[15] Committee ME-062, "Australian Standard - Use of ventilation and air conditioning in buildings (AS 1668.1:2015)," Standards Australia Limited., Standard, 2015.

[16] D. Dieziger, "Saving Money by Understanding Demand Charges on Your Electric Bill," http://www.fs.fed.us/t-d/pubs/htmlpubs/htm00712373/, 2000, [Online; accessed 21-01-2020].

[17] P. Hietaharju *et al.*, "A dynamic model for indoor temperature prediction in buildings," *Energies 2018*, vol. 11, no. 6, p. 1477, Jun 2018.

[18] G. Kenny *et al.*, "Towards establishing evidence-based guidelines on maximum indoor temperatures during hot weather in temperate continental climates," *Temperature*, vol. 6, no. 1, pp. 11–36, May 2019.

[19] N. R. E. Laboratory, "EnergyPlus," https://energyplus.net/, 2019, [Online; accessed 21-01-2020].

[20] C. Liu *et al.*, "Emergency Evacuation Model and Algorithm in the Building with Several Exits," *Procedia Engineering*, vol. 135, pp. 12–18, Dec 2016.

[21] R. Melfi *et al.*, "Measuring building occupancy using existing network infrastructure," in *Proc. IEEE IGCC and Workshops*, Orlando, USA, Jul 2011.

[22] M.J. Dixon, "Fire response of HVAC systems in multistory buildings: an examination of the nzb acceptable solutions," University of Canterbury, Research Report, 1998.

[23] I. Mohottige *et al.*, "Evaluating Emergency Evacuations Using Building WiFi Data," in *Proc. IEEE/ACM IoTDI*, Sydney, Apr 2020.

[24] National Fire Protection Association, "NFPA 101, Life Safety Code." https://www.nfpa.org/codes-and-standards, 2018, [Online; accessed 27-04-2019].

[25] S. news, "Sydney workers evacuated as smoke triggers alarms at offices, train stations," https://www.sbs.com.au/news/sydney-workers-evacuated-as-smoke-triggers-alarms-at-offices-train-stations, 2019, [Online; accessed 24-01-2020].

[26] L. Ng, "Changi Control Tower Evacuated Due To Fire Alarm," https://mustsharenews.com/changi-control-tower-evacuated/, 2020, [Online; accessed 24-01-2020].

[27] NSW Department of Planning, Industry and Environment, "Air Quality Data Services," https://www.dpie.nsw.gov.au/air-quality/air-quality-data-services/data-download-facility, 2020, [Online; accessed 13-01-2020].

[28] Occupational Safety and Health Administration, "Training Requirements in OSHA Standards," https://www.osha.gov/Publications, 2015, [Online; accessed 27-04-2019].

[29] Office of Environment and Heritage, "Optimising your heating, ventilation and air conditioning systems," NSW Gov., HVAC Optimisation Guideline, 2015.

[30] L. Perez-Lombard *et al.*, "A Review on Buildings Energy Consumption Information," *Energy and Build.*, vol. 40, no. 3, pp. 394–398, Jan 2008.

[31] Postlethwaite, S. and others, "Victorian Government Office Building Standards Guidelines," Victorian Government Property Group, Standard, 2001.

[32] S. Royer *et al.*, "Black-box modeling of buildings thermal behavior using system identification," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 10 850–10 855, Apr 2014.

[33] A. Trivedi *et al.*, "iSchedule: Campus-scale HVAC scheduling via mobile WiFi monitoring," in *Proc. ACM e-Energy*, Shatin, Hong Kong, May 2017.

[34] UNSW's Division of External Relations, "Annual Report," UNSW Sydney, Fact Sheet, 2019.

[35] A. Vishwanath *et al.*, "A Data Driven Pre-Cooling Framework for Energy Cost Optimization in Commercial Buildings," in *Proc. ACM e-Energy*, Shatin, Hong Kong, May 2017.

[36] G. Wang *et al.*, "Air handling unit supply air temperature optimal control during economizer cycles," *Energy and Build.*, vol. 49, pp. 310–316, Jun 2012.
[37] S. Zhang *et al.*, "Thermal comfort analysis based on PMV/PPD in cabins of manned submersibles," *Building and Environment*, vol. 148, pp. 668 – 676, Jan 2019.
[38] Z. Zheng *et al.*, "An edge based data-driven chiller sequencing framework for hvac electricity consumption reduction in commercial buildings," *IEEE Transactions on Sustainable Computing*, Jul 2019, early access.
[39] Zimmerman, G., "Old buildings become fire hazards without proper maintenance," https://www.facilitiesnet.com/firesafety/tip/Old-Buildings-Become-Fire-Hazards-Without-Proper-Maintenance--44311, 2019, [Online; accessed 24-04-2020].

Iresha Pasquel Mohottige received her B.Sc. degree in Electronic and Telecommunications Engineering from University of Moratuwa in Sri Lanka in 2014 and her Ph.D. in Electrical Engineering and Telecommunications from University of New South Wales (UNSW) in Sydney, Australia in 2021. Her primary research interests include data driven modeling, applied machine learning and computer networks.

Hassan Habibi Gharakheili received his B.Sc. and M.Sc. degrees of Electrical Engineering from the Sharif University of Technology in Tehran, Iran in 2001 and 2004 respectively, and his Ph.D. in Electrical Engineering and Telecommunications from the University of New South Wales (UNSW) in Sydney, Australia in 2015. He is currently a Senior Lecturer at UNSW Sydney. His research interests include programmable networks, learning-based networked systems, and data analytics in computer systems.

Arun Vishwanath (SM '15, M '11) is a lead research scientist at IBM Research in Melbourne, Australia working in the area of IoT for energy optimization in smart buildings and IoT security. He received the Ph.D. degree in Electrical Engineering from the University of New South Wales, Sydney, Australia in 2011 and was a visiting Ph.D. scholar in the Department of Computer Science, North Carolina State University, USA in 2008. His research interests span the areas of IoT applications, cybersecurity and software defined networking. Arun has received several awards from IBM for outstanding technical accomplishments. He is the recipient of the Best Paper Award at the ACM e-Energy 2018 conference, is appointed Distinguished Speaker of ACM and is a Senior Member of IEEE.

Salil S. Kanhere received his Ph.D. from Drexel University. He is a Professor in the School of Computer Science and Engineering at UNSW. His research interests include the Internet of Things, pervasive computing, blockchain, crowdsourcing, sensor networks, and security. He has published 200 peer-reviewed articles and delivered over 30 tutorials and keynote talks. He is a Senior Member of ACM. He is a recipient of the Humboldt Research Fellowship.

Vijay Sivaraman received his B. Tech. from the Indian Institute of Technology in Delhi, India, in 1994, his M.S. from North Carolina State University in 1996, and his Ph.D. from the University of California at Los Angeles in 2000. He has worked at Bell-Labs as a student Fellow, in a silicon valley start-up manufacturing optical switch-routers, and as a Senior Research Engineer at the CSIRO in Australia. He is now a Professor at the University of New South Wales in Sydney, Australia. His research interests include Software Defined Networking, network architectures, and cyber-security particularly for IoT networks.