Online HVAC-Aware Occupancy Scheduling with Adaptive Temperature Control

BoonPing Lim, Hassan Hijazi, Sylvie Thiébaux, and Menkes van den Briel

The Australian National University and NICTA, Canberra, Australia

Abstract. Heating, ventilation and air-conditioning (HVAC) is the largest consumer of electricity in commercial buildings. Consumption is impacted by group activities (e.g. meetings, lectures) and can be reduced by scheduling these activities at times and locations that minimize HVAC utilization. However, this needs to preserve occupants' thermal comfort and be responsive to dynamic information such as new activity requests and weather updates. This paper presents an *online* HVAC-aware occupancy scheduling approach which models and solves a joint HVAC control and occupancy scheduling problem. Our online algorithm greedily commits to the best schedule for the latest activity requests and notifies the occupants immediately, but revises the entire future HVAC control strategy each time it considers new requests and weather updates. In our experiments, the quality of the solution obtained by this approach is within 1% of that of the clairvoyant solution. We incorporate *adaptive comfort temperature control* into our model, encouraging energy saving behaviors by allowing the occupants to indicate their thermal comfort flexibility. In our experiments, the integration of adaptive temperature control further generates up to 12% of energy savings when a reasonable thermal comfort flexibility is provided.

1 Introduction

Heating, ventilation and air-conditioning (HVAC) dominates the energy consumption of commercial buildings, accounting for roughly 40% of the total building electricity consumption per annum [11,28]. With rising energy costs and increasingly stringent regulatory environments, improving the energy efficiency of HVAC operations in buildings has become an important issue.

Recent studies show that energy-oriented occupancy scheduling can lead to significant savings in energy consumption [5,19,20,21,22,23,24,26,27]. The idea is to *proactively* control occupancy in commercial offices and university buildings by scheduling energy-hungry activities such as meetings, workshops, lectures and exams, at times and locations that are favorable from an energy standpoint. Lim et al.'s "HVAC-aware" occupancy scheduling approach implements this idea by solving the *joint* HVAC control and occupancy scheduling problem [21,22], which consists in simultaneously optimizing the times and locations of the various activities and the HVAC control parameters at each time and building zone. By exploiting the synergy between HVAC control and occupancy scheduling, this approach achieves a much higher rate of energy savings than works that are based on (data-driven) black-box models of the HVAC control [5,20] or that minimize energy consumption proxies (e.g. number of rooms used) [24,27].

Unfortunately, with few exceptions [20], previous works focus on off-line scheduling, and assume that all activities to schedule and other parameters such as the weather forecast are known in advance. Existing approaches also adopt a fixed comfort temperature control, keeping the allowable temperature of occupied locations strictly within narrow bounds (e.g. 21°-23°C). Although both settings generate energy-efficient schedules, they nevertheless limit the practicability of the models in the real-world, and prevent further energy savings that could be obtained with more flexible temperature bounds. This paper presents two novel contributions which address these shortcomings. First, we extend Lim et al.'s HVAC-aware occupancy scheduling approach to process activity requests in an *on-line* manner. Second, we encourage energy saving behavior by allowing the occupant to indicate their thermal comfort flexibility, and use it in a principled way to introduce *adaptive comfort temperature control* in our model.

In more detail, we propose an on-line approach that models and solves the joint HVAC control and occupancy scheduling problem. Our on-line algorithm greedily optimizes (and commits to) the times and locations for the latest requests, leaving the rest of the future schedule fixed but revising the entire future HVAC control strategy. This ensures that whilst participants are instantly notified of the scheduled time and location for their requested activity, the HVAC control is constantly re-optimized and adjusted to the full schedule and weather updates. Our experiments demonstrate that the quality of the online solution is, on average, within 1% of that of the solution returned by the original clairvoyant HVAC-aware algorithm [22].

Adaptive comfort temperature control shifts away from fixed indoor comfort bands towards wider temperature operating bands. Recent work [1] shows that even a narrow variation of comfort temperatures can achieve significant energy savings. We introduce the notion of thermal comfort flexibility in our model by allowing occupants to indicate their level of tolerance to temperature fluctuation in the form of a) a threshold limiting the probability of temperature violation, and b) the maximum deviation allowed at any time. We then solve a robust optimization model which provide these probabilistic guarantees. Our experiments show that, when occupants are reasonably flexible, the integration of adaptive temperature control generates up an extra 12% of energy savings.

As an additional advantage, adaptive temperature control reduces the constrainedness of our online scheduling and control problem. This can make the problem solvable whereas fixed temperature bounds cannot be met, which often occurs for instance when a late request needs to be scheduled in the immediate future in a room whose current temperature is far away from the comfort band. In our experiments, adaptive temperature control solves 73% of the 700 instances that are unsolvable under the fixed temperature control model.

To summarize, the main contributions of the paper are: a) an efficient online model for the joint HVAC control and occupancy scheduling problem, b) a new notion of thermal comfort flexibility in energy-aware scheduling, c) experiments showing substantial energy reduction and improvement of solution feasibility over the state-of-the-art.

2 Related Work

Our work differs from previous work given its focus on: (i) a joint HVAC control and occupancy scheduling model which handles dynamically arriving scheduling requests, and (ii) an adaptive temperature control approach that allows the occupant to specify their thermal comfort flexibility. Existing works on energy-aware occupancy scheduling [5,21,22,23,24,26,27] focus on offline scheduling, and assume fixed comfort temperature setpoints. In reality, scheduling requests can arrive at any time of the day using existing room booking systems. A recent survey shows that 56% of meeting requests were made within 1 day before the actual meeting day [20]. Thus, the ability to handle impromptu requests is crucial. Moreover, the ability to update HVAC control following a change in forecast is also essential.

Kwak et al. [19,20] propose an online stochastic MILP model to schedule meetings. Their work calculates energy consumption based on historical data and exploits flexibility in the time and location at which a meeting can take place. However, it does not optimize HVAC control, nor does it take thermal comfort flexibility into account. Our results show that combining meeting scheduling with HVAC control, and enabling adaptive temperature control based on occupant thermal comfort flexibility, significantly impacts energy savings.

Conventionally, room temperature is maintained within strict comfort bounds while occupied. Such control is not the most effective, since the HVAC system tries to achieve fixed temperature setpoints regardless of ambient conditions or the comfort levels of the individual occupants. More recent work enables adaptive thermal comfort control [1,6,7,18,25,31,32], exploiting the observation that when occupants have some form of input to the control, their subjective view of comfort changes and they are more willing to accept wider operating conditions than those mandated by traditional comfort models. For example, when controlling the emissivity of dynamic windows to reduce HVAC consumption in a smart home, Ono et al. [25] allow for temperature bound violations, but limit their probability using chance constraints with occupant-specified thresholds. Inspired by these works, we introduce the notion of thermal comfort flexibility into our scheduling model. We incorporate occupants' tolerance level as an input, allowing the scheduler to identify the best location and time slots that optimize energy saving while satisfying occupant thermal comfort.

Energy-oriented scheduling has gained more attention in recent years due to the significant cost saving opportunities. Ifrim et al. [17] present a MIP-based energy-price savings scheduling model to reduce cost in production scheduling. Dupont et al. [8] use CP to develop an energy aware framework for virtual machine placement in cloud-based data centers. Scott et al. [29] describe an online stochastic MILP to schedule home appliances based on real-time pricing. Most works focus on energy-aware scheduling in production lines, data centers and residential buildings whilst our work specifically targets energy-efficient scheduling in the smart building space, which is dominated by HVAC consumption.

3 Online Occupancy Scheduling

This section presents our online occupancy scheduling problem. We start by describing the scheduling setting and our notations. We then cover the scheduling constraints and variables which, later on in Section 4, will interact with the HVAC control model to form a more complex joint scheduling and control model. We formulate our model as a mixed-integer program (MIP). It can be solved using a MIP solver, or when scaling up to problems of practical size, by combining MIP with large neighborhood search (LNS) as explained in [22].

In our online setting, the scheduler runs recurrently and each run is called an *online session*. Each online session $i \in I$ starts at time τ_i and ends before the next session starts at time τ_{i+1}. The scheduling and control model discretizes time into a set K of *time steps*. Each time step $k \in K$ starts at time t_k. Two consecutive time steps k and $k+1$ are separated by a fixed duration $t_{k+1} - t_k = \boldsymbol{\Delta_t} \in \mathbb{R}^+$. Each on-line session i considers a horizon of $\boldsymbol{n}$ time steps $K(i) = \{k(i), \ldots, k(i)+\boldsymbol{n}-1\}$ where $k(i)$, the first time step in that horizon, is the least time step in K such that $t_{k(i)} \geq \tau_i$.

Let L be the set of locations (or, interchangeably, zones) in the building, and P be a set of participants. An activity request m is a tuple $\langle \boldsymbol{a}_m, K_m, L_m, P_m, \boldsymbol{d}_m, F_m \rangle$ where $\boldsymbol{a}_m \in \mathbb{R}^+$ is the request arrival time, $K_m \subseteq K$ is the set of time steps at which the activity is permitted to <u>start</u> in the future (for each $k \in K_m, \boldsymbol{a}_m < t_k$), $L_m \subseteq L$ is the set of locations at which the activity is permitted to take place, $P_m \subseteq P$ is the set of attendees for the activity, $\boldsymbol{d}_m \in \mathbb{N}$ is the activity duration (number of time steps), and F_m represents the comfort temperature flexibility parameters which will be explained in Section 4.4. Note that the sets K_m and L_m can be used to encode a variety of situations, such as room capacity requirements, availability of special equipment such as video conferencing, time deadlines for the activity, and attendee availability constraints. We write $\mathcal{C}(M)$ for the set of attendee conflicts w.r.t. a set of requests M; each conflict C is a subset of requests, each pair of which has at least one attendee in common: $\mathcal{C}(M) = \{C \subseteq M \mid \forall m, m' \in C, P_m \cap P_{m'} \neq \emptyset\}$.

To account for all activities that have been scheduled so far, we maintain a master schedule S as a set of triples $\langle m, l, k \rangle$ storing the activity request id m, the assigned location l, and the time step k at which m is scheduled to start. At each online session i, the scheduler schedules the new activity requests $N(i)$ which have been received since the start of session $i-1$, i.e, each $m \in N(i)$ satisfies $\tau_{i-1} < \boldsymbol{a}_m \leq \tau_i$. It also needs to consider, without modifying them, the set $Q(i)$ of ongoing activities and future activities that were scheduled during previous sessions: $Q(i) = \{m \mid \exists \langle m, l, k \rangle \in S \text{ such that } k + \boldsymbol{d}_m - 1 \geq k(i)\}$. So overall, the set of activities to consider at session i is $M(i) = N(i) \cup Q(i)$. To simplify the scheduling model below, we assume that for each pre-scheduled request $m \in Q(i)$

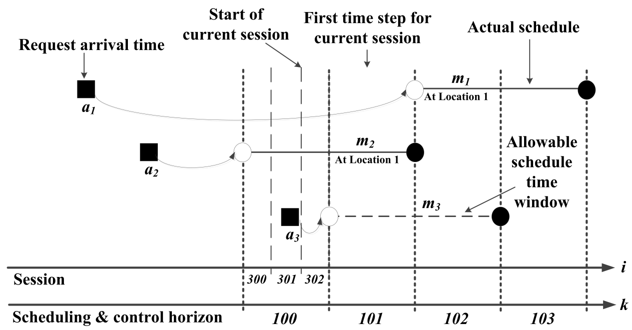

Fig. 1: Online scenario

such that $\langle m, l, k\rangle \in S$, the set of permissible locations is reduced to $L_m = \{l\}$, and the set of permissible start time steps is reduced to the scheduled start time k or the first time step $k(i)$ of the session, which ever occurs last, i.e. $K_m = \{\max(k, k(i))\}$. For consistency, the meeting duration $\boldsymbol{d}_m$ is decremented by $k(i) - k$; this is only needed later in Section 4.4 for equation (23).

Fig. 1 shows a scenario example featuring three requests m_1, m_2 and m_3 with arrival times $\boldsymbol{a}_1$, $\boldsymbol{a}_2$ and $\boldsymbol{a}_3$, respectively. The set of locations is $L = \{\boldsymbol{l}_1, \boldsymbol{l}_2\}$. The dash vertical lines show the start of the sessions, and the dotted vertical lines delimit the time steps. In this instance, the scheduler runs every 10 minutes and the time steps are 30 minutes long. At the start of session $i = 302$, requests m_1 and m_2 have already been scheduled and m_3 is a new request, hence $N(302) = \{m_3\}$, $Q(302) = \{m_1, m_2\}$. The master schedule is $S = \{\langle m_1, \boldsymbol{l}_1, 102\rangle, \langle m_2, \boldsymbol{l}_1, 100\rangle\}$, and the first time step of the new session is $k(302) = 101$. The set of permissible locations and start time steps for the new request are $K_3 = \{101, 102\}$ and $L_3 = \{\boldsymbol{l}_1, \boldsymbol{l}_2\}$ ($\boldsymbol{l}_1$ will be ruled out by the scheduler). Those of the pre-existing requests are reduced as follows: $L_1 = \{\boldsymbol{l}_1\}, K_1 = \{102\}, L_2 = \{\boldsymbol{l}_1\}$ and $K_2 = \{101\}$.

We are now ready to describe our scheduling constraints and variables for online session i. The main scheduling variable is the boolean decision variable $x_{m,l,k}$ which is true iff request $m \in M(i)$ is scheduled to take place at zone $l \in L_m$ starting at time slot $k \in K_m$. We also introduce the variables $y_{m,l,k}$ which is true iff activity m is scheduled to occupy location l at time step k, $z_{l,k}$ which is true iff zone l is occupied at time step k, and $pp_{l,k}$ which indicates the number of people in zone l at time step k. These variables will be used by the HVAC control part of the model in 4.

The scheduling constraints are the following. Constraints (1) ensure that all requests are scheduled exactly once within the allowable start times and locations. Constraints (2) define the $y_{m,l,k}$ variables. Constraints (3) state that no more than one activity can occupy a location at any time and define the $z_{l,k}$ variables. Observe that the right hand side of these constraints is either zero or one, which limits the number of activities to at most one. Also, when the left hand side equals one then the zone must be occupied. Constraints (4) determine the number $pp_{l,k}$ of occupants at each location and time step, and finally constraints (5) ensure that activities with at least one attendee in common cannot be scheduled in parallel. Once a new request $m \in N(i)$ has been scheduled, the master schedule S is updated by adding the 3-tuple $\langle m, l, k \rangle$ for which $x_{m,l,k} = 1$.

$$\sum_{l \in L_m, k \in K_m} x_{m,l,k} = 1 \quad \forall m \in M(i) \tag{1}$$

$$\sum_{\substack{l \in L_m, \; k' \in K_m: \\ k - \boldsymbol{d}_m + 1 \leq k' \leq k}} x_{m,l,k'} = y_{m,l,k} \quad \forall m \in M(i), l \in L, k \in K(i) \tag{2}$$

$$\sum_{m \in M(i)} y_{m,l,k} \leq z_{l,k} \quad \forall l \in L, k \in K(i) \tag{3}$$

$$\sum_{m \in M(i)} y_{m,l,k} \times |P_m| = pp_{l,k} \quad \forall l \in L, k \in K(i) \tag{4}$$

$$\sum_{m \in \nu, l \in L_m} y_{m,l,k} \leq 1 \quad \forall k \in K(i), \nu \in \mathcal{C}(M(i)) \tag{5}$$

4 HVAC Control Model

This section covers the HVAC control model and the adaptive temperature control approach. We describe the HVAC control aspects starting with the objective function we consider, the effect of the control on the building thermal dynamics, and the fixed temperature bounds – we refer the reader to [21] for a more detailed treatment. We subsequently extend the model with adaptive temperature control to further maximize energy savings.

4.1 Variable-Air-Volume Systems

Following Goyal et al. [14,15], we focus on commercial buildings with variable-air-volume (VAV) based HVAC systems, which serve over 30% of the commercial building floor space in the United States [9]. A schematic of a VAV-based HVAC system with two VAV boxes connected to two building zones is shown in Fig. 2.

The air handling unit (AHU) supplies conditioned air to the VAV boxes. The AHU consumes energy when mixing outdoor air with return air and cooling it to the pre-set conditioned air temperature T^{CA} [12.8 °C]; it consumes less energy when the outdoor air temperature T^{OA} is closer to T^{CA}. Each VAV box consumes energy when regulating the supply air temperature T^{SA} and the supply air flow rate a^{SA} to keep the zone temperature T within comfort bounds; in particular, it may need to reheat the conditioned air. Finally, the supply fan at the AHU consumes energy to maintain a constant air pressure through the supply duct; it may speed up or slow down depending on air flow rates used by the VAV boxes.

We focus on control strategies that can be applied to each VAV box. For such strategies, the key HVAC decision variables are the supply air flow rate $a^{SA}_{l,k}$ and

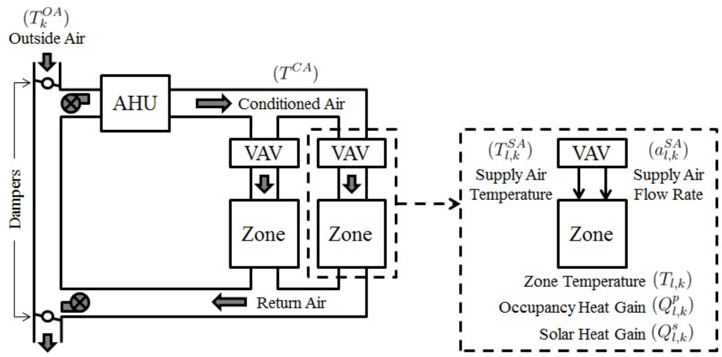

Fig. 2: VAV-based HVAC system.

temperature $T_{l,k}^{SA}$ at each zone/location $l \in L$ and time step $k \in K$. We determine an optimal control for these variables, given the occupancy schedule and the bounds on supply air temperature, supply air flow rate, and room temperature during vacant and occupied periods.

4.2 Objective Function

Specifically, our goal is to generate energy-efficient schedules that minimize the energy use of air-conditioning, re-heating and fan operations of the HVAC. Thus, the objective function for online session i is the following.

$$\text{minimize} \sum_{k \in K(i)} \left(p_k^{cond} + p_k^{fan} + \sum_{l \in L} p_{l,k}^{heat} \right) \times \boldsymbol{\Delta_t} \tag{6}$$

where

$$p_k^{cond} = \boldsymbol{C}^{pa} \left(T_k^{OA}(i) - \boldsymbol{T}^{CA} \right) \sum_{l \in L} a_{l,k}^{SA} \quad \forall k \in K(i) \tag{7}$$

$$p_k^{fan} = \boldsymbol{\beta} \sum_{l \in L} a_{l,k}^{SA} \quad \forall k \in K(i) \tag{8}$$

$$p_{l,k}^{heat} = \boldsymbol{C}^{pa} (T_{l,k}^{SA} - \boldsymbol{T}^{CA}) a_{l,k}^{SA} \quad \forall l \in L, k \in K(i) \tag{9}$$

Constraints (7)-(9) determine the values of the variables p_k^{cond}, p_k^{fan}, $p_{l,k}^{heat}$, which respectively represent the energy consumed by the AHU for conditioning, by the supply fan for maintaining air pressure, and by the VAV box for reheating the conditioned air. In constraint (7), we assume that online session i uses the latest update $T_k^{OA}(i)$ available for the outdoor temperature forecast at each time step k. The coefficients in these constraints are the fan power coefficient $\boldsymbol{\beta}$ (0.65), and the heat capacity of air at constant pressure $\boldsymbol{C}^{pa}$ (1.005 kJ/kg·K).

$$\boldsymbol{T}^{CA} \leq T_{l,k}^{SA} \leq \overline{\boldsymbol{T}}^{SA} \quad \forall l \in L, k \in K(i) \tag{10}$$

$$\underline{\boldsymbol{a}}^{SA} \leq a_{l,k}^{SA} \leq \overline{\boldsymbol{a}}^{SA} \quad \forall l \in L, k \in K(i) \tag{11}$$

Moreover, constraints (10) and (11) ensure that the supply air temperature and the air flow rate are bounded by the HVAC operational capacity. The supply air temperature $T_{l,k}^{SA}$ may range from that of the conditioned air $\boldsymbol{T}^{CA}$ (12.8 °C), up to $\bar{\boldsymbol{T}}^{SA}$ (40 °C) if the air is reheated at the VAV box. The air flow rate $a_{l,k}^{SA}$ can fluctuate between $\underline{\boldsymbol{a}}^{SA}$ (0.108 kg/s) and $\bar{\boldsymbol{a}}^{SA}$ (5.0 kg/s), where the lower bound is determined by the ASHRAE ventilation standard and the upper bound is reached when the VAV dampers are fully open.

4.3 Building Thermal Dynamics

Next, we want our control to appropriately constrain zone temperatures. The first step to do this is to introduce a new variable $T_{l,k}$ representing the temperature at each zone and time step, and model the effects of the HVAC control on this zone temperature. To capture the building thermal dynamics, we adopt a computationally efficient lumped RC-network [12,13,14] which incorporates the thermal resistance and capacitance of each zone and between adjacent zones, the latest available forecast of the solar gain $Q_{l,k}^{s}(i)$, and the internal heat gain $Q_{l,k}^{p}$ generated by the occupants at each zone. The latter is directly proportional to the number of occupants $pp_{l,k}$ scheduled to be at the zone by the online scheduler – this is one of the variables via which the scheduling and control models interact. We use a discrete-time linear model

$$T_{l,k+1} = f_l(T_{l,k}, u_{l,k}, v_{l,k}) \quad \forall l \in L, k \in K(i) \tag{12}$$

where $u_{l,k} = [a_{l,k}^{SA}, T_{l,k}^{SA}, pp_{l,k}]$ is the vector of controllable variables, and $v_{l,k} = [Q_{l,k}^{s}(i), T_{k}^{OA}(i)]$ is the vector of exogenous inputs. With this model, the HVAC control is optimized over the entire horizon $K(i)$. E.g., the optimal control could activate the HVAC at night to benefit from the low outside night temperature to pre-cool a room for an early morning meeting. See [21] for details.[1]

4.4 Adaptive Temperature Control

Having modeled the effect of the HVAC control on the zone temperatures $T_{l,k}$, we are now ready to ensure that the HVAC fulfills its main role of keeping these zone temperatures within appropriate comfort bounds. In the fixed comfort bound model found in much of the literature, when a zone is occupied, the zone temperature must lie within a specified comfort interval $[\underline{\boldsymbol{T}}, \bar{\boldsymbol{T}}]$ ([21 °C, 23 °C]). When the zone is empty, its temperature can fluctuate more freely within $[\underline{\boldsymbol{T}}^{\emptyset}, \bar{\boldsymbol{T}}^{\emptyset}]$ (16 °C, 28 °C]). These bounds can be set to reflect individual building guidelines. As shown in [21], maintaining temperature within these fixed bounds can be achieved by adding constraints (13) to our model. In these constraints,

[1] Both Lim et al. [21,22] and our experiments use a more complex state vector which not only includes the zone temperatures $T_{l,k}$ but also the temperature of the interior walls. For readability reasons, we abstract from these extra state variables in our exposition above.

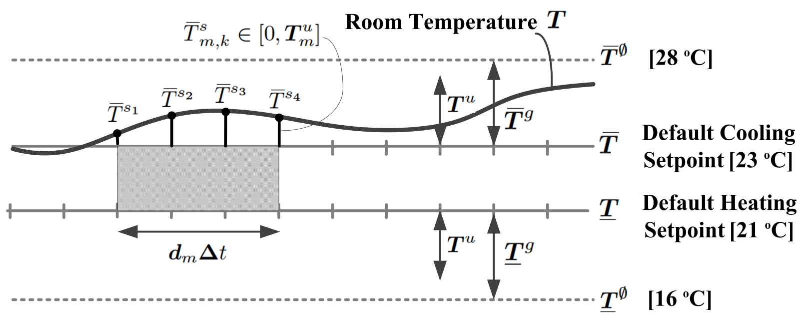

Fig. 3: Adaptive Temperature Control

the HVAC model interacts with the scheduling model via the variables $z_{l,k}$ that indicate whether or not location l is occupied at time step k. The constants $\underline{\boldsymbol{T}}^g$ and $\overline{\boldsymbol{T}}^g$ denote the gap between the occupied and unoccupied temperature lower and upper bounds.

$$\underline{\boldsymbol{T}}^{\emptyset} + \underline{\boldsymbol{T}}^g z_{l,k} \le T_{l,k} \le \overline{\boldsymbol{T}}^{\emptyset} - \overline{\boldsymbol{T}}^g z_{l,k} \quad \forall l \in L, k \in K(i) \tag{13}$$

In the present paper, we generate additional energy savings by departing from these fixed comfort bounds. We adopt a flexible temperature bound model, in which the comfort interval is dynamically configured through input parameters reflecting the flexibility of occupants. Specifically, the input parameters we consider for an activity request m are $F_m = \langle \boldsymbol{T}^u_m, \boldsymbol{\alpha}_m, \boldsymbol{p}_m \rangle$ and are such that the HVAC control will guarantee: a) that the zone temperature will never exceed $[\underline{\boldsymbol{T}} - \boldsymbol{T}^u_m, \overline{\boldsymbol{T}} + \boldsymbol{T}^u_m]$ at any point during the activity and b) that with probability at least $\boldsymbol{p}_m$, the cumulative temperature violation during the activity will be bounded by $\boldsymbol{\alpha}_m$. The parameter $\boldsymbol{\alpha}_m$ is equivalent to the duration for which the occupant would be willing to let the temperature deviation be $\boldsymbol{T}^u_m$. Fig. 3 illustrates these concepts. In this example, activity m occupies location l for 3 times steps. The occupant is prepared to accept a maximal deviation (of up to 3°) from the default comfort bounds (i.e. [18 °C, 26 °C]), but also wants the cumulative violation to remain within acceptable bounds (the equivalent of 20 min at 3°) with high probability (0.9). This is achieved by setting $\boldsymbol{T}^u_m = 3$, $\boldsymbol{\alpha}_m = 20$, and $\boldsymbol{p}_m = 0.9$.

Let m be a meeting scheduled to start at time step $j \in K_m$ in location l. To formalize these concepts, we introduce the following slack variables in the model $\underline{T}^s_{m,k} \in [0, \boldsymbol{T}^u_m]$ and $\overline{T}^s_{m,k} \in [0, \boldsymbol{T}^u_m]$, for $k \in K(i)$. These variables represent our unknown temperature violations above and below the default bounds $[\underline{\boldsymbol{T}}, \overline{\boldsymbol{T}}]$. Based on these variables, the first guarantee we want to provide can be written as the adaptive counterpart of the fixed temperature bound constraints (13).

$$\underline{\boldsymbol{T}}^{\emptyset} + \underline{\boldsymbol{T}}^g z_{l,k} - \underline{T}^s_{m,k} \le T_{l,k} \le \overline{\boldsymbol{T}}^{\emptyset} - \overline{\boldsymbol{T}}^g z_{l,k} + \overline{T}^s_{m,k} \tag{14}$$

The second guarantee is about bounding the cumulative temperature violation, this can be formulated as follows,

$$\sum_{k=j}^{j+\boldsymbol{d}_m-1} \left(\underline{T}^s_{m,k} + \overline{T}^s_{m,k}\right) \boldsymbol{\Delta}_t \le \boldsymbol{\alpha}_m \boldsymbol{T}^u_m \tag{15}$$

To implement a probabilistic version of this constraint, we introduce independent uniformly distributed random variables $\rho_{m,k} \in [-1, 1]$, which represent the noise in our temperature violation, transforming constraints (15) into,

$$\sum_{k=j}^{j+\boldsymbol{d}_m-1} \left(\underline{T}^s_{m,k} + \overline{T}^s_{m,k} - \rho_{m,k} \right) \boldsymbol{\Delta}_t \leq \boldsymbol{\alpha}_m \boldsymbol{T}^u_m \tag{16}$$

We then resort to results from the Robust Optimization literature [10,3,4,2,16] to be able to offer the following probabilistic guarantee,

$$Pr\left(\sum_{k=j}^{j+\boldsymbol{d}_m-1} \left(\underline{T}^s_{m,k} + \overline{T}^s_{m,k} - \rho_{m,k} \right) \boldsymbol{\Delta}_t \leq \boldsymbol{\alpha}_m \boldsymbol{T}^u_m \right) \geq \boldsymbol{p}_m \tag{17}$$

where $Pr\left(f_{\boldsymbol{\rho}}(x) \leq 0\right)$ denotes the probability of satisfying constraint $f_{\boldsymbol{\rho}}(x) \leq 0$ given the uncertainty created by the random variables $\boldsymbol{\rho}$. In particular, based on [2, Theorem 3.], we can offer the above probabilistic guarantee by enforcing the following constraint

$$\sum_{k=j}^{j+\boldsymbol{d}_m-1} \rho^2_{m,k} \leq \boldsymbol{\delta}^2_m, \tag{18}$$

where the the ellipsoid radius $\boldsymbol{\delta}_m$ is linked to the constraint satisfaction probability $\boldsymbol{p}_m$ as follows:

$$\boldsymbol{p}_m \geq 1 - exp(-\boldsymbol{\delta}^2_m/1.5).$$

For instance, a radius of $\boldsymbol{\delta}_m = 2.63$ leads to a constraint satisfaction probability $\boldsymbol{p}_m \geq 0.99$. Furthermore, based on [16, Corollary 1.], we can write the following deterministic equivalent of (17) without having to explicitly enforce (18),

$$\sum_{k=j}^{j+\boldsymbol{d}_m-1} \left(\underline{T}^s_{m,k} + \overline{T}^s_{m,k} \right) - |\mathcal{S}| - \sqrt{\left(\boldsymbol{\delta}^2_m - |\mathcal{S}|\right) |\overline{\mathcal{S}}|} \leq \boldsymbol{\alpha}_m \boldsymbol{T}^u_m / \boldsymbol{\Delta}_t, \tag{19}$$

where the set $\mathcal{S}$ is described in [16, prop 1.]. For computational efficiency reasons, this is the approach we adopt in our current implementation.

Since activity locations and start times are not known in advance, we introduce variables $\underline{T}^{\xi}_{l,k}$ (resp. $\overline{T}^{\xi}_{l,k}$) such that $\underline{T}^{\xi}_{l,k} = \underline{T}^s_{m,k}$ and $\overline{T}^{\xi}_{l,k} = \overline{T}^s_{m,k}$ when activity $m \in M(i)$ occupies location $l \in L$ at time slot $k \in K(i)$, i.e., when $y_{m,l,k} = 1$. In order to accommodate activities that span multiple scheduling horizons, we also introduce the inputs $\boldsymbol{T}^{prev}_m = \sum_{k \in K: k < k(i)} (\underline{\boldsymbol{T}}^s_{m,k} + \overline{\boldsymbol{T}}^s_{m,k})$, which accounts for the amount of cumulative violation consumed before the start of the current session. Recall also from Section 3 that meetings that have been scheduled in previous sessions have their start time set $\boldsymbol{K}_m$, location set $\boldsymbol{L}_m$ and duration $\boldsymbol{d}_m$ reduced accordingly when the current session starts. With these notations, the overall adaptive temperature control constraints replacing the fixed temperature constraints (13) in the HVAC control model are the following.

$$\underline{\boldsymbol{T}}^{\emptyset} + \underline{\boldsymbol{T}}^{g} z_{l,k} - \underline{T}^{\xi}_{l,k} \leq T_{l,k} \leq \overline{\boldsymbol{T}}^{\emptyset} - \overline{\boldsymbol{T}}^{g} z_{l,k} + \overline{T}^{\xi}_{l,k} \quad \forall l \in L, k \in K(i) \tag{20}$$

$$\underline{T}^{s}_{m,k} - \hat{\boldsymbol{T}}(1 - y_{m,l,k}) \leq \underline{T}^{\xi}_{l,k} \leq \underline{T}^{s}_{m,k} + \hat{\boldsymbol{T}}(1 - y_{m,l,k}) \ \forall m \in M(i), l \in L, k \in K(i) \tag{21}$$

$$\overline{T}^{s}_{m,k} - \hat{\boldsymbol{T}}(1 - y_{m,l,k}) \leq \overline{T}^{\xi}_{l,k} \leq \overline{T}^{s}_{m,k} + \hat{\boldsymbol{T}}(1 - y_{m,l,k}) \ \forall m \in M(i), l \in L, k \in K(i) \tag{22}$$

$$\sum_{k=j}^{j+\boldsymbol{d}_m-1} \left(\underline{T}^{s}_{m,k} + \overline{T}^{s}_{m,k}\right) - |\mathcal{S}| - \sqrt{(\boldsymbol{\delta}^2_m - |\mathcal{S}|)\,|\overline{\mathcal{S}}|} \leq \boldsymbol{\alpha}_m \boldsymbol{T}^u_m / \boldsymbol{\Delta}_t - \boldsymbol{T}^{prev}_m, \forall m \in M(i), j \in \boldsymbol{K}_m \tag{23}$$

$$\underline{T}^{\xi}_{l,k} \leq \sum_{m \in M(i)} \boldsymbol{T}^u_m y_{m,l,k} \quad l \in L, k \in K(i) \tag{24}$$

$$\overline{T}^{\xi}_{l,k} \leq \sum_{m \in M(i)} \boldsymbol{T}^u_m y_{m,l,k} \quad l \in L, k \in K(i) \tag{25}$$

Constraints (20) are the adaptive bound constraints. Constraints (21-22) are the on-off constraints defining the variables $\underline{T}^{\xi}_{l,k}$ and $\overline{T}^{\xi}_{l,k}$ with $\hat{\boldsymbol{T}} = \max_{m \in M(i)} \{\boldsymbol{T}^u_m\}$. Constraint (23) is the probabilistic constraint on the cumulative temperature violation, taking into account $\boldsymbol{T}^{prev}_m$. The last two constraints force the corresponding slack to zero when a location is unoccupied.

5 Experimental Results

5.1 Problem Sets

We analyze our contributions using 9 problem sets with increasing numbers of activities (meetings) and locations (meeting rooms). The problem sets are labeled 10M-4R, 20M-20R, 50M-20R, 100M-20R, 200M-20R, 50M-50R, 100M-50R, 200M-50R, and 500M-50R, where xM-yR consists of problem instances with x meetings and y rooms. Each set contains 800 problem instances, giving a total of 7200 instances, obtained as follows.

We start from a set of real data from 32,065 unique meetings in a USC library collected by Kwak [20]. Each meeting request in this original data set includes the request arrival time, start time, duration, specified room and number of attendees. We first derive a probability distribution on meeting start times from this data set. To obtain a set of requests, we sample x meetings for this distribution. We then create different instances with that set of requests by varying the time flexibility, request-to-start time gap, and temperature flexibility of the requests. The time flexibility of a request m is its number $|K_m| \in \{1, 2, 4, 8, 32\}$ of permissible start time steps. The request-to-start time gap denotes the duration {10 minutes, 1 hours, 4 hours, 24 hours} between the request's arrival time $\boldsymbol{a}_m$ and its first possible start time step. The temperature flexibility indicates the level of tolerance for the room temperature deviation from the standard heating (21°C) and cooling (23°C) setpoints, and is one of three settings: low, medium, or high flexibility, with $\boldsymbol{p}_m = 0.99$ for all settings, $\boldsymbol{T}^u_m = 2$ (low), 3 (medium), 5 (high), and $\boldsymbol{\alpha}_m = 10$ (low), 20 (medium), 30 (high).

Note that in the high setting, the deviation could be up to 5°C, which is equivalent to 28°C for 30 minutes. This is an extreme case used to study the effects of temperature flexibility, but not a recommended setting. In the more realistic medium setting, the deviation is only up to 3°C, which is equivalent to 26°C for 20 minutes.

We keep the meeting duration and number of attendees identical to that of the original meeting request from the USC data, and assume that the occupant is fully flexible in terms of location, that is, that the meeting can be allocated to any room. In all problem sets, the duration $\mathbf{d}_m$ of meetings ranges from 1 to 4 time steps (30 minutes to 2 hours). The meetings must be scheduled over a period of 5 summer days. The available rooms are located in 5 buildings and differ by their thermal resistance and capacitance [22]. We use a 1×4 zone layout where each zone has the same thermal resistance and capacitance as its neighboring zones. Moreover, all rooms have the same geometric area of $6 \times 10 \times 3$ m^3 with a window surface area of 4×2 m^2 and a capacity of 30 people. The solar gain ranges from 50 to 350 W/m^2 during the day. All activities have between 2 and 30 attendees. All our experiments were run on a cluster consisting of a 2 $\times$ AMD 6-Core Opteron 4334, 3.1GHz with 64GB memory.

5.2 Solution Method

To solve these problem instances, we combine our MIP model with Large Neighborhood Search as explained in [22]. LNS is a local search metaheuristic which iteratively improves an initial solution by alternating between a destroy and a repair step [30]. In brief, our LNS approach works as follows.

In every online session i, we start by generating an initial feasible solution, in two steps. First, we find a feasible occupancy schedule that minimizes the number of rooms used. Second, we determine the HVAC control settings (supply air flow rate and temperature) that minimize energy consumption for this schedule.

Our destroy step destroys part of the schedule by unscheduling the subset of new requests $N(i)$ that are allocated to two to four randomly selected locations. This forms an energy-aware meeting scheduling subproblem that is much smaller than the original problem and can be solved effectively using MIP. The repair step consists in repairing the schedule and re-optimizing the entire HVAC control by solving this subproblem using our MIP model. If this leads to an improved solution, then the new schedule and control settings are accepted. Otherwise, we keep the solution that was just destroyed. Given that the LNS starts with a feasible solution and does not accept infeasible solutions, the solution remains feasible throughout the execution of the algorithm.

5.3 Online vs. Offline Scheduling

We start by comparing the solution quality of our online approach with that of the offline approach [22]. In the online approach, the scheduler runs LNS for 5 minutes in each session, with a MIP runtime limit of 8 seconds in each iteration. In the offline approach, the entire set of requests to schedule is given,

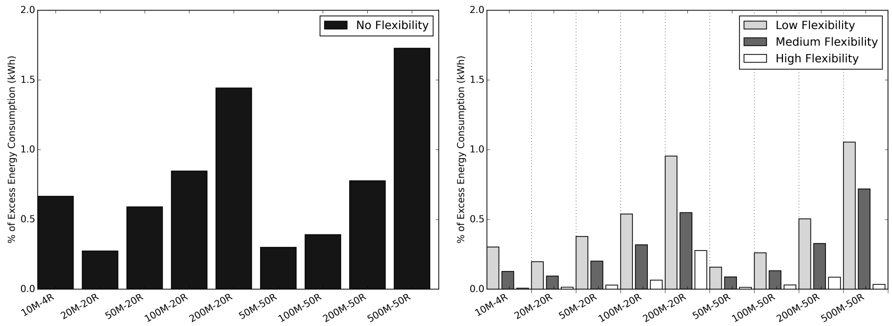

Fig. 4: Online vs. Offline Scheduling: With Fixed Temperature Setpoints (left) and Adaptive Temperature Setpoints (right)

and we compute the final schedule; The scheduler runs LNS for 2 hours, with a MIP runtime limit of 15 minutes in each iteration. To identify how much more improvement can be obtained, we warm start the offline schedule with the best online solution found (over all the possible request-to-start time gaps).

The difference of solution quality, that is the excess consumption of the online scheduling as a percentage of the off-line scheduling consumption, is shown in Fig. 4. The results for the fixed temperature setpoints are shown on the left, whilst those for the adaptive temperature setpoints are on the right. Both graphs show that the offline solutions are merely 1% to 1.5% better than the online solutions for tightly constrained problems (such as 200M-20R, 500M-50R), and that, as expected, the online approach improves when the problem is less constrained in terms of meetings to rooms ratio and temperature flexibility. Note that in the online approach, at most 20 requests arrive in each online session and a maximum of 4 rooms are destroyed, thus the sub-problems formed are small enough for MIP to solve them to (near) optimality. The off-line approach has many more meetings to deal with, but on the other hand, as problems become more constrained, it has more room to optimize than the greedy on-line approach. Altogether, even with a simple greedy approach, our online algorithm is able to perform effectively without prior knowledge of future requests.

5.4 Energy Savings of Adaptive Temperature Control

Next, we examine the benefits of our adaptive temperature control, which allows the occupant to specify their level of tolerance for the room temperature deviation from the fixed 21°-23°C comfort bounds. Because HVAC consumption is highly dependent on the temperature gap between the outdoor temperature and the occupied temperature setpoint, we show that even a small variation from the original setpoints can lead to large energy savings.

Fig. 5 shows the additional energy savings obtained with adaptive temperature control as a percentage of the fixed temperature control consumption (left), and the maximum temperature deviation incurred by the adaptive approach (right). The left figure shows that the additional savings can reach up to [8%, 12.7%, 16.5%] depending on the [low, medium, high] temperature flexibility allowed by the occupants. The right figure shows that the maximum degree of

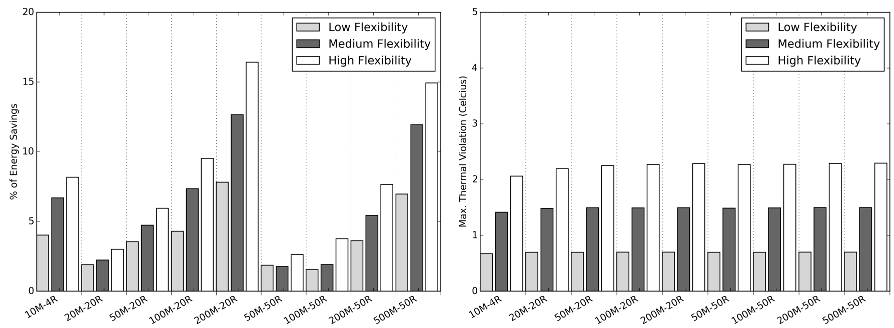

Fig. 5: Energy Savings from Adaptive Temperature Control (left) and Maximum Temperature Deviation from Standard Setpoints ($21° - 23°$C)

temperature deviation is only about [0.7, 1.5, 2.3]°C for low-to-high temperature flexibility, respectively. Overall, increasing temperature flexibility reduces HVAC consumption and cost. Taking an energy rate of \$0.24/kWh and the 500M-50R problem set as example, this corresponds to annual savings of about [\$11500, \$19542, \$24690] for [low, medium, high] temperature flexibility.

5.5 Model Feasibility

Finally, we study the solution feasibility of on-line scheduling with fixed and adaptive temperature control, respectively. Fig. 6 shows the percentage of feasible solutions generated by the two approaches, as a function of the request-to-start time gap. Altogether, adaptive temperature control solves 73% of the instances that are deemed unsolvable under the fixed temperature control regime.

We observed that with fixed temperature setpoints, we fail to generate feasible solutions in most cases when the requests arrive less than 1 hour prior to the earliest possible activity start time. This infeasibility issue mainly happens at the initialization stage, where the initial schedule generation is decoupled from the initial HVAC control generation. In order to quickly generate an initial feasible schedule, activities are packed into the minimum number of rooms possible. However, the room temperatures may be too far from the temperature setpoints to obtain an initial feasible HVAC control reaching the designated occupied temperature at short notice. In contrast, the model with adaptive temperature control is able to solve many of these problem instances, and even generates some feasible solutions when the requests arrive just 10 minutes prior to the earliest activity start time. This is mainly due to the relaxation of the temperature setpoints. We observed that the number of feasible solutions increases proportionally to the temperature flexibility.

Apart from the constrainedness imposed on temperature setpoints, the model also stumbles into infeasibility when the scheduler fails to schedule all requests due to the lack of feasible location or time slot. Overall, the performance improves as the request-to-start time gap increases for both models.

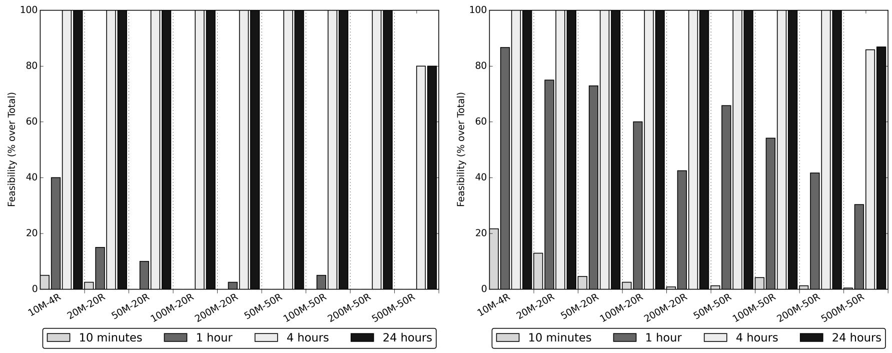

Fig. 6: Solution Feasibility: With Fixed Temperature Setpoints (left) and Adaptive Temperature Setpoints (right)

6 Conclusions & Future Work

In this paper we develop an online scheduling model and adaptive temperature control method for joint HVAC control and occupancy scheduling. Leveraging an explicit model of building occupancy-based HVAC control, our model adopts a greedy approach to schedule dynamically arriving requests to take place at locations and times that are favorable from energy standpoint. Our experiments show that, even without prior knowledge of future requests, our model is able to produce energy-efficient schedules which are less than 1% away from the clairvoyant solution.

We extend the model to enable adaptive temperature control, moving away from the conventional fixed comfort temperature setting. The occupant is allowed to indicate their level of tolerance for the room temperature to deviate from the standard heating and cooling setpoints. We shows that thermal comfort flexibility significantly impacts energy consumption. Compared to the existing fixed temperature control, the energy savings in our experiments can reach up to 8% with low temperature flexibility, with a maximum deviation of 0.7°C from the original setpoints, and up to 15% with high temperature flexibility with a maximum of 2.3°C deviation from the standard setpoints. We have also shown that given some thermal comfort flexibility, our model is able to schedule requests arriving 10 minutes prior to the start time, and produce substantially more feasible solutions than the conventional fixed temperature setpoints approach.

We are interested in exploring new algorithmic approaches that allows us to improve our solution and scale even further. We are particularly interested in investigating stochastic scheduling and control, which allows us to predict future request arrival and cancellations. We are also interested in exploring the CP formulation of joint HVAC control and meeting scheduling. As the joint model consists of hybrid discrete-continuous variables, we plan to reformulate it by discretizing the HVAC control variables, and compare the solution quality generated by both MIP and CP models.

Acknowledgements

Thanks to Milind Tambe and Jun Kwak from USC for sharing the real data in [20] and helpful discussions. This work is supported by NICTA's Optimization Research Group as part of the Future Energy Systems project. NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

References

1. Aileen, E.: The potential energy savings through the use of adaptive comfort cooling setpoints in fully air conditioned australian office buildings, a simulation study. Equilibrium Journal (2010)
2. Babonneau, F., Vial, J.P., Apparigliato, R.: Uncertainty and Environmental Decision Making. International Series in Operations Research and Management Science, Springer-Verlag (2009)
3. Ben-Tal, A., Nemirovski, A.: Robust convex optimization. Mathematics of operations research 23(4), 769–805 (1998)
4. Ben-Tal, A., Nemirovski, A.: Robust solutions to uncertain linear programs. OR Lett. 25, 1–13 (1999)
5. Chai, B., Costa, A., Ahipasaoglu, S.D., Huang, S., Yuen, C., Yang, Z.: Minimizing commercial building cost in smart grid: An optimal meeting scheduling approach. In: Smart Grid Communications (SmartGridComm), 2014 IEEE International Conference on. pp. 764–769. IEEE (2014)
6. Chew, B., Kazi, S., Amiri, A.: Adaptive thermal comfort model for air-conditioned lecture halls in malaysia. World Academy of Science, Engineering and Technology, International Journal of Civil, Environmental, Structural, Construction and Architectural Engineering 9(2), 150–157 (2015)
7. De Dear, R.J., Brager, G.S., Reardon, J., Nicol, F., et al.: Developing an adaptive model of thermal comfort and preference/discussion. ASHRAE transactions 104, 145 (1998)
8. Dupont, C., Giuliani, G., Hermenier, F., Schulze, T., Somov, A.: An energy aware framework for virtual machine placement in cloud federated data centres. In: Future Energy Systems: Where Energy, Computing and Communication Meet (e-Energy), 2012 Third International Conference on. pp. 1–10. IEEE (2012)
9. EIA: Us eia-department of energy, cbecs detailed tables (2003). (2003), http://www.eia.gov/consumption/commercial/
10. El Ghaoui, L., Lebret, H.: Robust solutions to least-squares problems with uncertain data. SIAM Journal on Matrix Analysis and Applications 18, 1035–1064 (1997)
11. of Energy, D.: Buildings energy data book. In: Buildings Energy Data Book. Department of Energy, United States (2011), http://buildingsdatabook.eren.doe.gov/ChapterIntro1.aspx
12. Gouda, M., Danaher, S., Underwood, C.: Low-order model for the simulation of a building and its heating system. Building Services Engineering Research and Technology 21(3), 199–208 (2000)
13. Gouda, M., Danaher, S., Underwood, C.: Building thermal model reduction using nonlinear constrained optimization. Building and Environment 37(12), 1255–1265 (2002)

14. Goyal, S., Barooah, P.: A method for model-reduction of non-linear thermal dynamics of multi-zone buildings. Energy and Buildings 47, 332–340 (2012)
15. Goyal, S., Ingley, H.A., Barooah, P.: Occupancy-based zone-climate control for energy-efficient buildings: Complexity vs. performance. Applied Energy 106, 209–221 (2013)
16. Hijazi, H., Bonami, P., Ouorou, A.: Robust delay-constrained routing in telecommunications. Annals of Operations Research 206(1), 163–181 (2013)
17. Ifrim, G., OSullivan, B., Simonis, H.: Properties of energy-price forecasts for scheduling. In: Principles and Practice of Constraint Programming. pp. 957–972. Springer (2012)
18. Klein, L., Kwak, J.y., Kavulya, G., Jazizadeh, F., Becerik-Gerber, B., Varakantham, P., Tambe, M.: Coordinating occupant behavior for building energy and comfort management using multi-agent systems. Automation in Construction 22, 525–536 (2012)
19. Kwak, J.y., Kar, D., Haskell, W., Varakantham, P., Tambe, M.: Building thinc: User incentivization and meeting rescheduling for energy savings. In: Proc. 13th International Conference on Autonomous Agents and Multi-agent Systems. pp. 925–932 (2014)
20. Kwak, J.y., Varakantham, P., Maheswaran, R., Chang, Y.H., Tambe, M., Becerik-Gerber, B., Wood, W.: Tesla: An energy-saving agent that leverages schedule flexibility. In: Proc. 12th International Conference on Autonomous Agents and Multi-agent Systems. pp. 965–972 (2013)
21. Lim, B.P., Van Den Briel, M., Thiébaux, S., Backhaus, S., Bent, R.: Hvac-aware occupancy scheduling. In: AAAI. pp. 4249–4250 (2015)
22. Lim, B.P., Van Den Briel, M., Thiébaux, S., Bent, R., Backhaus, S.: Large neighborhood search for energy aware meeting scheduling in smart buildings. In: Integration of AI and OR Techniques in Constraint Programming, pp. 240–254. Springer (2015)
23. Majumdar, A., Zhang, Z., Albonesi, D.: Characterizing the benefits and limitations of smart building meeting room scheduling. In: Proc. 7th International Conference on Cyber-Physical Systems (2016)
24. Majumdar, A., Albonesi, D.H., Bose, P.: Energy-aware meeting scheduling algorithms for smart buildings. In: Proc. 4th ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings. pp. 161–168. ACM (2012)
25. Ono, M., Graybill, W., Williams, B.C.: Risk-sensitive plan execution for connected sustainable home. In: Proc. 4th ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings. pp. 45–52. ACM (2012)
26. Pan, D., Wang, D., Cao, J., Peng, Y., Peng, X.: Minimizing building electricity costs in a dynamic power market: Algorithms and impact on energy conservation. In: Real-Time Systems Symposium (RTSS), 2013 IEEE 34th. pp. 107–117. IEEE (2013)
27. Pan, D., Yuan, Y., Wang, D., Xu, X., Peng, Y., Peng, X., Wan, P.J.: Thermal inertia: Towards an energy conservation room management system. In: Proc. 31st IEEE International Conference on Computer Communications. pp. 2606–2610. IEEE (2012)
28. Pitt, Sherry (eds.): Baseline Energy Consumption and Greenhouse Gas Emissions in Commercial Buildings in Australia Part 1 Report. Department of Climate Change and Energy Efficiency, Australia (2012), http://www.industry.gov.au/Energy/EnergyEfficiency/Non-residentialBuildings/Documents/CBBS-Part-1.pdf

29. Scott, P., Thiébaux, S., Van Den Briel, M., Van Hentenryck, P.: Residential demand response under uncertainty. In: Principles and Practice of Constraint Programming. pp. 645–660. Springer (2013)
30. Shaw, P.: Using constraint programming and local search methods to solve vehicle routing problems. In: Proc. 4th International Conference on Principles and Practice of Constraint Programming (CP). pp. 417–431 (1998)
31. Ward, J., Wall, J., White, S.: Automate and motivate: behaviour-reliant building technology solutions for reducing greenhouse gas emissions. Architectural Science Review 53(1), 87–94 (2010)
32. Yang, R., Wang, L.: Development of multi-agent system for building energy and comfort management based on occupant behaviors. Energy and Buildings 56, 1 – 7 (2013)