Coordinated Scheduling of Thermostatically Controlled Real-Time Systems under Peak Power Constraint

Gopinath Karmakar[1], Ashutosh Kabra
Bhabha Atomic Research Centre
Mumbai, India

Krithi Ramamritham[2]
Indian Institute of Technology, Bombay
Mumbai, India

Abstract—**A substantial fraction of the energy demand of buildings comes from air-conditioners (ACs), refrigerators, etc., which do not need human interaction for their continuous operations. So long as desirable temperature levels, which we refer to as *thermal comfort-bands*, are maintained by such *background* loads, users will not be concerned about *when* they perform their assigned functions. i.e., *when* they consume the energy required to function. This paper addresses the problem of maintaining thermal comfort-bands associated with background loads under Peak energy consumption constraint. Based on our empirical observations pertaining to energy consumption profiles of thermal comfort-band maintaining ACs, a) we present a feasibility criterion for maintaining the thermal comfort-band of a given set of background loads under constraints on the peak available power and (b) we propose TCBM, a new algorithm for scheduling such background loads under peak power constraint. In addition to limiting peak power demand, the TCBM algorithm avoids undesirable switching (ON and OFF) of electrical appliances to improve efficiency of the equipment and reduce failures. Results from simulation and real-life implementation demonstrate that our algorithm is superior to the existing work on background load scheduling. We also show how TCBM can adapt to changes in ambient parameters and provide the basis for efficient demand-response systems.**

Index Terms—**TCBM; Smart Home; Green home; Real-Time**

I. INTRODUCTION

A significant fraction of home as well as large buildings' energy demand comes from air-conditioners, refrigerators and room-heaters. These Thermostatically Controlled Electrical Devices (TCEDs or TCE devices) maintain the temperature of the environment under their control according to user-specified set-points. They do not run continuously, but follow a pattern of ON-OFF cycles to maintain the desired temperature. The TCE devices, which we also refer to as *background loads*, do not need human interaction for their continuous operation. So long as desirable temperature levels, which we refer to as *thermal comfort-bands*, are maintained by such *background* loads, users will not be concerned about *when* they perform their assigned functions, i.e., when they consume the energy required to perform these functions.

Under un-coordinated individual control, it is possible that all the above devices run simultaneously during some time

intervals. This can result in higher peak power demand. Today many commercial buildings are subject to very high tariff under peak power demand pricing [1]. In the near future, residential buildings are also likely to be subject to such electricity pricing schemes. Further, predictably lower peak demand benefits power utilities by reducing the cost and complexity of handling power grid stability, load-sheding and blackouts. Hence, the design and evaluation of techniques for the coordinated scheduling of *background* loads form the crux of this paper. The objective of any TCE device, A_i, is to maintain the temperature T_i of the environment under its control within a desirable thermal *comfort-band* $[T^U, T^L]$. Under peak power constraint, not all devices may be able to run concurrently at any point of time. Therefore, under peak power demand constraint, we need to switch power between the devices intelligently, while ensuring that the desired temperature is maintained by each. Suppose W_i denotes the wattage of device A_i and PPL denotes the Peak Power Limit, then the devices that can be powered at any point of time will be governed by the following constraints:

$$\forall t, \sum_{i=1}^{n} x_i(t) \times W_i \leq PPL \qquad (1)$$

where, $\forall i, x_i(t) \in \{0, 1\}$ and $T^L \leq T_i \leq T^U$. $x_i(t)$ is the state (1 = ON and 0 = OFF) of the i^{th} device at time t.

We focus on the background load due to ACs, as analysis of an AC system will be applicable to other TCE devices, like refrigerators, dehumidifier, room heaters, etc. All these devices maintain the comfort-band of the environment that they control and work on the same basic principles of heat transfer.

The functioning of a TCE device A_i can be modelled as a periodic activity having an ON-time and OFF-time within its *duty cycle* equivalent to execution time (C_i) and laxity (L_i) respectively; the duty cycle is hence equivalent to period $P_i = C_i + L_i$. For an AC unit that maintains temperature T_i within a comfort band $[T^U, T^L]$, we can consider C_i as the time duration the device runs to bring down T_i from T^U to T^L and L_i as the time duration it can be switched OFF, i.e., when T_i rises from T^L to T^U. Any *individual* TCE device exhibits this periodic nature when controlled thermostatically.

Not surprisingly, existing literature [2] [3][4] suggests modelling background electrical loads as real-time tasks and
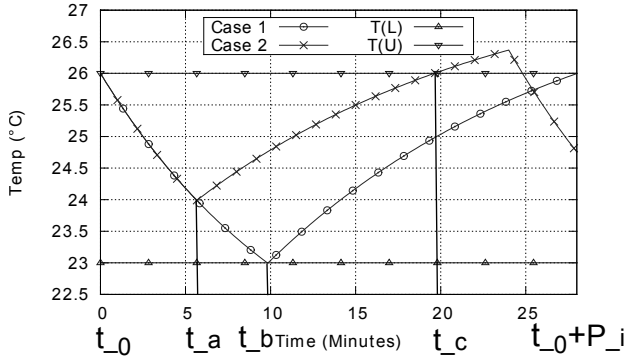
Fig. 1. TCED charasteristics

applying traditional real-time scheduling algorithms EDF[5], LSF[6], etc. The main assumption in real-time schedulability analysis is that a task is considered schedulable, if it is executed for C_i (here, the duration for which the TCED is ON, i.e., receives power) units of time every P_i period and it receives C_i units of processor time before it deadline $D_i(C_i \leq D_i \leq P_i)$; *when* this C_i time is received within P_i along the time scale is not a consideration.

We now show that in order to schedule thermostatically controlled electrical devices, it is not just enough to allocate power for C_i units of time within P_i; it is also important *when* this C_i is allocated within P_i. It can be seen in Figure 1, that if the TCED task is executed (powered-ON) for C_i units of time during the beginning of its period from t_0 to t_b, the environmental temperature T_i remains within the comfort-band. But, if the task is preempted after its execution till t_a and is powered-ON for the remaining execution time just at the end of the period $[t_i, t_0 + P_i]$, T_i goes beyond the upper limit T^U of the comfort-band $[T^U, T^L]$.

Thus, whereas traditional real-time scheduling algorithms consider the deadline D_i, maximum laxity L_i and duty cycle P_i as constants for a task, in the case of TCE devices these parameters are dynamic, because execution of such devices depend on the existing temperature of the very environment controlled by them. From Figure 1, it can be observed that for the same TCED task, the duty cycle should change dynamically from P_i to $(t_c - t_0)$ when preempted at $t = t_a$. Also, the maximum Laxity L_i changes from $(P_i - t_b)$ to $(t_c - t_a)$. Therefore, in case of TCED tasks these parameters are required to be calculated dynamically in order to apply any scheduling algorithm effectively.

Furthermore, preemption decisions must be taken based on the environmental parameter, which is to be maintained within the comfort-band $[T^U, T^L]$. For example, so as not to violate the need for $T_i \leq T^U$ of AC A_i, another AC A_j, whose zonal temperature T_j is $< T^U$ might have to be preempted to divert power to A_i.

Preemption decisions must also consider the fact that minimal switching of TCE devices, except for resistive heating loads, is desirable for the following reasons:

1) Compressor driven devices have a specified delay (~ 3

min) before they can be restarted. The delay allows the pressures in the system to equalize so that the compressor does not start under a load. If no restart-delay is provided, the compressor may not start due to an overload or it can even damage the equipment.

2) High starting current [7] with every switch-ON of induction motors (used in ACs, refrigerators, etc.) causes additional power loss.

These considerations motivated us to develop a new algorithm, Thermal Comfort-Band Maintenance (TCBM), to schedule thermostatically controlled devices (presented in Section III). The basic intuition behind TCBM is that since the main goal is to maintain the temperature within a desirable $[High, Low]$ band, we need not do switching (ON/OFF) of any load, if the temperature is within the comfort-band associated with it. This helps in reducing the number of undesirable switching of devices. TCBM is also designed to support feasibility analysis, presented in Section III-D, using which it will be possible to determine, at the time a user sets the desired comfort-band, whether that comfort band can be maintained given peak power constraint.

Besides TCBM and the associated feasibility analysis, this paper makes the following additional contributions.

- A conceptual model of TCE devices as a special class of real-time tasks (in Section II). The model is corroborated up by our empirical study of the thermal characteristics of ACs, which demonstrates that the temperature of the environment controlled by an AC rises exponentially, when the device is switched OFF and falls exponentially when it is switched ON. The exponential nature of the characteristic equations underlies the design of the TCBM algorithm as well as the analysis of feasibility of maintaining the specified comfort-band under peak demand constraint.

- Refining the concepts of Laxity (L_i), Deadline (D_i) and Period (P_i) to make them cognizant of the environmental parameter (temperature T_i) controlled by a TCE device. This sets the stage for adapting scheduling algorithms from the real-time literature for TCED scheduling.

- A performance study, involving simulation as well as prototype implementation (presented in Section IV) that demonstrates the superior performance of our algorithm compared to approaches adapted from the literature.

- Algorithms for on-line *adaptive determination of comfort-bands* of TCE devices based on our feasibility analysis and insights obtained from experimental observations (discussed in Section V).

Related work is discussed in Section VI and is followed by conclusion. Important parameters and notations used in the paper are listed in Table I.

II. DETERMINING THE THERMAL CHARACTERISTICS OF ACS

We consider a typical office building or home having AC units cooling thermally de-coupled zones. A zone can be a

TABLE I
PARAMETERS & NOTATIONS

Parameter	Notation
Total no. of Devices	n
No. of Devices that can run at a time	m
Power Requirement of i^{th} AC	W_i
Scheduling Decision Interval	I^S
Temperature in i^{th} Zone/Room	T_i
Ambient Temperature	T_a
Warming slope when i^{th} AC OFF	S_r^i
Cooling slope when i^{th} AC ON	S_f^i
Upper value of zonal temperature	T^U
Lower value of zonal temperature	T^L
Min. T_i to switch-ON a device	$T^U - \Delta^U$
Desired T_i to switch-OFF a device	$T^L + \Delta^L$

TABLE II
CONSTANTS FOR AC (ON) CHARACTERISTIC EQUATION

Ambient temperature = $27.5^0 C$		
Constants	**AC1**	**AC2**
a	2.765	1.281
b	0.008	0.01336
c	0.6791	1.602
d	22.5	20.49

room with a single AC or a large room is assumed to be divided into thermally de-coupled zones whose temperatures are maintained by one AC in each zone. We have a set of n ACs for n zones, but peak power constraint permits at most m ACs out of n to be ON at any point of time. The ACs are to be scheduled such that the temperature $T_i, i = 1, 2, \cdots n$ in individual zones lies within the comfort-band $[T^U, T^L]$.

As portrayed in Fig. 1, when an AC is on, i.e., the compressor is running, it cools down the room temperature T_i; the temperature rises from the instant the AC is switched OFF because of heat loads and losses. Based on elementary principles of heat transfer [8], we develop a model, simplified by means of using overall heat transfer co-efficient in calculating heat transfer from the terminal temperatures, i.e. the temperature of the two bodies between which heat transfer takes place. It may be noted that various heat transfer co-efficients, pertaining to different heat transfer modes (conduction, convection and radiation), are combined into an overall heat transfer co-efficient for simplification of the heat transfer problem as is done in practice [9].

Using this (simplified) model we first derive the warming and cooling rate of an AC. The time taken by an AC to cool a zone from T_i to T^L and the time it takes for T_i to rise up to T^U, when the AC is OFF are then derived, based on experimental data.

A. Modeling Rate of Change in Temperature

- *Warming when AC is OFF:* Let Q_{hi} be the heat input rate in zone Z_i, $1 \le i \le n$ and T_i be the temperature of the zone. Then, the rate of change of T_i can be expressed as

$$S_i \frac{dT_i}{dt} = Q_{hi} + h_0(T_a - T_i) \qquad (2)$$

where S_i denotes the thermal capacity of the zone and h_0 denotes the overall heat transfer co-efficient of the room. It follows that

$$\frac{dT_i}{dt} = \frac{Q_{hi} + h_0 T_a}{S_i} - \frac{h_0}{S_i} T_i = \alpha' - \beta' T_i \qquad (3)$$

in which $\alpha' = \frac{Q_{hi} + h_0 T_a}{S_i}$ and $\beta' = \frac{h_0}{S_i}$.

- *Cooling when AC is ON:* When an AC is switched ON, it removes heat from the room. With heat load Q_{hi} and

ambient temperature T_a remaining constant, Equation 2 can be modified to capture the heat removal by AC as

$$S_i \frac{dT_i}{dt} = Q_{hi} - h_1(T_i - T^C) + h_0(T_a - T_i) \qquad (4)$$

where, T^C is the temperature of the heat-transfer coil of the AC and h_1 is the overall heat transfer co-efficient of the AC. It follows that

$$\frac{dT_i}{dt} = \frac{Q_{hi} + h_1 T^C + h_0 T_a}{S_i} - \frac{h_1 + h_0}{S_i} T_i = \alpha - \beta T_i \qquad (5)$$

in which $\alpha = \frac{Q_{hi} + h_1 T^C + h_0 T_a}{S_i}$ and $\beta = \frac{h_1 + h_0}{S_i}$.

B. Experimentally Determining AC Thermal Characteristics

We measured the changing temperatures of a zone controlled by an AC using $Pt100$ RTD (Resistance Temperature Detectors) and recorded them every 10 sec. in a digital recorder (Eurotherm Chessell 5000). The recorder provides a facility to generate alarm(s), when the input parameter(s) goes beyond specified limit(s). We generated cooling curve data by switching ON an AC when the room temperature T_i was equal to the ambient temperature $T_a = 27.5^0 C$ and kept it running till T_i saturated to the lowest value ($20.5^0 C$ here). Drop in temperature gets bounded because, at some point, the rate of cooling matches with the rate of heating due to heat loads and losses. The data of warming curve was generated by switching OFF the AC at $T_i = 20.5^0 C$ and keeping it off till T_i went up to $27.5^0 C$ From the experimental data, we did curve-fitting and obtained the thermal characteristic equations of ACs:

- *Cooling Down (AC Switched ON):*

$$T_i(t) = ae^{-bt+c} + d \qquad (6)$$

where, a, b, c & d are constants specific to a particular AC for a particular ambient temperature, the values of which obtained by curve-fitting from our experimental data for two ACs are shown in Table II.

Solving Equation 6, we get

$$t = \frac{1}{b}(c - ln\frac{T_i - d}{a}) \qquad (7)$$

Therefore, the time for T_i to reach from T^U to T^L, i.e.,

$$C_i = \frac{1}{b}[(c - ln\frac{T^U - d}{a}) - (c - ln\frac{T^L - d}{a})] \qquad (8)$$

$$C_i = \frac{1}{b} ln\frac{T^L - d}{T^U - d} \qquad (9)$$

Also, from Equation 6 we get

$$\frac{dT_i}{dt} = -abe^{-bt+c} = -abe^{-bt+c} - bd + bd.$$

TABLE III
CONSTANTS FOR AC (OFF) CHARACTERISTIC EQUATION

Ambient temperature = $27.5^0 C$		
Constants	**AC1**	**AC2**
a'	-2	-2.391
b'	0.009	0.009317
c'	1	0.9455
d'	28	27.7

TABLE IV
AVERAGE C AND P OF ACs FOR VARYING COMFORT-BANDS

(a) With Equal Comfort-bands

Comfort-band (^0C)	$C_i(Min.)$	$P_i(Min.)$	C_i/P_i
$25.0 - 22.0$	14	28	0.5
$25.5 - 22.5$	12	28	0.429
$26.0 - 23.0$	10	29	0.345

(b) With Comfort-bands of varying T^U

Comfort-band (^0C)	$C_i(Min.)$	$P_i(Min.)$	C_i/P_i
$25.5 - 22.0$	16	34	0.471
$26.0 - 22.0$	17	40	0.425
$26.5 - 22.0$	18	46	0.391

(c) With Comfort-bands of varying T^L

Comfort-band (^0C)	$C_i(Min.)$	$P_i(Min.)$	C_i/P_i
$26.0 - 23.0$	10	29	0.345
$26.0 - 22.5$	14	34	0.412
$26.0 - 22.0$	18	40	0.450

Substituting the value of T_i from Equation 6, we get

$$\frac{dT_i}{dt} = bd - bT_i \tag{10}$$

- *Warming up (AC Switched OFF)*

$$T_i(t) = a'e^{-b't+c'} + d' \tag{11}$$

where a', b', c' & d' are constants specific to a particular AC at for a particular ambient temperature obtained by curve-fitting from our experimental data. Solving Equation 11, we get

$$t = \frac{1}{b'}[c' - ln\frac{T_i - d'}{a'}]. \tag{12}$$

Therefore, the time for T_i to reach from T^L to T^U, i.e.,

$$L_i = \frac{1}{b'}[c' - ln\frac{T^U - d'}{a'}] - \frac{1}{b'}[c' - ln\frac{T^L - d'}{a'}]. \tag{13}$$

$$L_i = \frac{1}{b'}ln\frac{T^U - d'}{T^L - d'}. \tag{14}$$

Also, from Equation 11 we get

$$\frac{dT_i}{dt} = -a'b'e^{-b't+c'} = -a'b'e^{-b't+c'} - b'd' + b'd'$$

Considering the value of T_i from Equation 11, we get

$$\frac{dT_i}{dt} = b'd' - b'T_i \tag{15}$$

It may be noted that Equations 10 and 15 obtained from experimental data have exactly the same form as that of Equations 5 and 3 respectively, obtained by our simplified theoretical model described in Section II-A. Hence, they corroborate each other.

C. Effect of Comfort-Band Settings on AC Operation

We operated an AC in an office room on a regular working day and studied the effect of comfort-band $[T^U, T^L]$ on the cooling time C_i and the period P_i of AC on three different comfort-bands as follows. (The results in Table IV are based on the average of the results obtained from 10 duty-cycles and rounded off to nearest integer.)

- *Comfort-Bands with equal* $[T^U - T^L]$*:* In this experiment, the position of the comfort-band is shifted along the temperature scale while keeping $[T^U - T^L]$ the same. It can be observed from the experimental data presented in Table IV(a) and from Figure 2(a) that the cooling time C_i as well as the *duty cycle* $(C_i + P_i)$ depend not only on the size of the comfort band but also on the value of T^L. It can also be observed from Figure 2(a), that shifting of comfort-band by just 0.5^0C can cause

a significant change in the cooling time C_i and in the utilization C_i/P_i.

These observations are in compliance with Eq. (9). The increase in C_i can be explained from the fact that $ln\ (A - d)/ln\ (B - d) > ln\ A/ln\ B$, if $A < B$, $0 < d < A$ and $0 < d < B$. The scenario is the same with the Eq. (9), when the comfort-band is changed from $[26.0^0C, 23.0^0C]$ to $[25.5.0^0C, 22.5^0C]$ or to $[25.0^0C, 22.0^0C]$.

- *Comfort-Bands with varying* T^U*:* We increased the T^U values by 0.5^0C, keeping the T^L values same, as shown in Table IV(b) and observed that by increasing T^U by 0.5^0C, the decrease in C_i/P_i is larger, as expected.
- *Comfort-Bands with varying* T^L*:* When the values of T^L were decreased by 0.5^0C, keeping the values of T^U the same, as shown in Table IV(c), we observed that decreasing the lower limit T^L of the comfort-band, as expected, increases the value of C_i/P_i.

From Figures 2 $(a), (b)$ & (c) and from Equation 10, it can be observed that the higher the temperature, the faster is the rate of cooling, because of the exponential nature of the cooling curve. Therefore, an AC will run for lesser time if T^U is shifted up as compared to the same when T^U is lower. This results in reduced power consumption.

We have shown here i) how to obtain the constants a, b, c, d and a', b', c', d' experimentally, to characterize the ACs ii) how manipulating the comfort-band can affect the AC power consumption. These observations are used in Section V for adaptive demand-response control.

III. TCBM SCHEDULING

The primary goal of TCBM is to maintain the comfort-band. TCBM achieves this while turning ON ACs judiciously because of the peak power constraint. Secondly, turning-ON an AC has to be done taking into account the fact that once turned-OFF, it can not be turned-ON for a minimum period of time, typically 3 min. Thirdly, the number of switchings (ON

a) Comfort-bands of equal width ($T^U - T^L$)

b) Comfort-bands with varying T^U


c) Comfort-bands with varying T^L

Fig. 2.   Effect of Comfort-Band on AC Operation

& OFF) is to be minimized for reasons given earlier. TCBM scheduling achieves this by allowing an AC to remain ON till its zonal temperature reaches T^L or it becomes necessary to switch it OFF in the event the zonal temperature of some other AC reaches T^U. A_i, once switched OFF, is kept OFF till T_i reaches T^U, preventing switching of ACs when all T_is are within the comfort-band.

In order to maintain zonal temperatures within comfort-band under the peak power demand constraint, before switching ON an AC which has reached T^U, another AC may have to be switched-OFF . Since one of our aims is to reduce switching of ACs to a minimum, it is logical to select the *coolest* [1] AC.

---

[1]*Coolest* AC is the one that takes maximum time to reach T^U from its current T_i.

The scheduling (ON-OFF) decisions are taken periodically. Therefore, if a switching decision is made only after the zonal temperature reaches T^U, comfort-band can be violated before the next scheduling decision is taken. To avoid the possibility of comfort-band violation, an AC is considered for switching ON or OFF ahead in time before its temperature reaches T^U or T^L. Switching ON of an AC may also be delayed if more than m ACs reach T^U simultaneously.

Definition 1: B^U is the upper-limit of T_i only above which, the i^{th} AC is considered for switching ON. The value of B^U is calculated as $B^U = T^U - \Delta^U$, where Δ^U denotes the rise in T_i due to the maximum possible delay in switching ON the AC after it reaches B^U.

Definition 2: B^L is the lower-limit of T_i below which, the i^{th} AC is switched OFF unless it becomes necessary to power on some other AC to maintain the comfort-band. The value of B^L is calculated as $B^L = T^L + \Delta^L$, where Δ^L denotes the fall in T_i due to the maximum possible delay in switching OFF the AC after it reaches B^L.

A. The TCBM Algorithm

Initially, $T_i = T_a$ and an AC operates to bring down the T_i from T_a to a value that lies within the comfort-band. Specifically, the TCBM algorithm starts by switching ON m arbitrarily chosen ACs, where the value of m is determined by the peak demand constraint, so that

$$\sum_{i=1}^{m} W_i \leq PPL \qquad (16)$$

where, W_i denotes the power required by i^{th} AC and the sum is obtained from the first m of the n ACs arranged in descending order of their wattage.

Once T_i is within the comfort band, the ACs are controlled to maintain the respective T_is within the comfort band, by applying the following rules every I^S units of time.

1) *Rule # 1:* Turn OFF AC_i if it is ON at time t and if
   a) $T_i \leq B^L$ *OR*
   b) there is an $AC_j (i \neq j)$ with $T_j \geq B^U$ *AND* no. of ON-ACs $\geq m$ *AND* $T_i \leq T^U$ *AND* AC_i is the *coolest* one among ON-ACs.

2) *Rule # 2:* Turn ON AC_i if
   a) $T_i \geq B^U$ *AND*
   b) No. of ON-ACs $< m$

B. Calculating Δ^U

Since an AC is switched ON only if its zonal temperature reaches B^U, the worst delay in switching-ON an AC will arise when all ACs are OFF and the zonal temperatures T_i of all the ACs reaches B^U at a time, say $t = t_s$.

Now, at t_s, only m ACs can be switched-ON and therefore, the T_i corresponding to $(n-m)$ ACs will rise beyond B^U. Our goal is to switch-ON these $n-m$ ACs within a reasonable time so that none of the T_i will cross T^U. Let I^S be the scheduling period. Therefore, after $t_s + I^S$ time, the next set of m ACs will be running, either by switching-ON ACs which were OFF

during the previous scheduling period or allowing ACs to run, which were already ON.

Therefore, the best case will be, if at $t = t_s + I^S$ the next set of m ACs are selected from the set of ACs which were OFF and in every subsequent scheduling period, a new set of m ACs are selected in the same manner till all the ACs are switched ON at least once after t_s. Thus the AC, which will be switched ON last after t_s, will face a delay t_d, where

$$t_d = \lceil \frac{n}{m} \rceil \times I^S \tag{17}$$

Now, by definition, Δ^U is the temperature rise from B^U to T^U. Therefore, switching ON of an AC, should be allowed at least Δ^U units of temperature ahead of T^U to ensure that the AC will not cross T^U because of the delay in scheduling decision, as discussed. Further, a minimum Δ^U ensures minimum number of AC switching.

From Equation 15, we get the temperature rising slope at B^U as

$$\frac{dT_i}{dt} \mid_{T_i = B^U} = b'(d' - B^U) \tag{18}$$

Substituting B^U with $T^U - \Delta^U$, [Definition 1]

$$\frac{dT_i}{dt} \mid_{T_i = B^U} = b'(d' - T^U + \Delta^U) \tag{19}$$

Considering the worst case of linear rise of temperature from B^U to T^U, we get

$$\frac{\Delta^U}{t_d} = b'd' - b'(T^U - \Delta^U)$$

Substituting the value of t_d from Equation 17, we get

$$\Delta^U = \frac{b'(d' - T^U)}{1 - b'\lceil \frac{n}{m} \rceil I^S} \times \lceil \frac{n}{m} \rceil I^S \tag{20}$$

C. Calculating Δ^L

Since any number of ACs can be turned OFF at an instance, there is no delay involved beyond I^S.

Considering the temperature falling slope according to Equation 10 and following a reasoning similar to that described in Section III-B, we get

$$\Delta^L = \frac{b(d - T^L)}{1 + bI^S} \times I^S \tag{21}$$

D. Feasibility Analysis

In this subsection, we discuss the feasibility of scheduling m out of n ACs that ensures a given comfort-band $[T^U, T^L]$.

Theorem 1: A given comfort-band $[T^U, T^L]$ will be maintained by scheduling m out of n ACs, if

$$\sum_{i=1}^{m} abs(b_i(d_i - B^U)) \geq \sum_{i=1}^{n-m} b'_i(d'_i - B^U) \tag{22}$$

Where, i) b_i, d_i and b'_i, d'_i are the constants pertaining to the cooling curve and the warming curve respectively corresponding to i^{th} AC, as discussed in Section II-B and ii) sum on the left is obtained from the first m of the n ACs arranged in ascending order of their falling slope and iii) sum on the right is obtained from the first $(m - n)$ of the n ACs arranged in descending order of their rising slope.

Proof Sketch: In Section III-B, it has been shown that, in order to ensure minimum switching, scheduling decision to switch ON an AC is taken when its zonal temperature reaches B^U. Therefore, we check for feasibility considering the worst case when, say at time t_0, the zonal temperatures (T_i) of all the ACs reach B^U because it is a scenario when all the n ACs are ready to be switched ON.

Under peak demand constraint, at time t_0, only m ACs will be switched ON and in every successive scheduling period, a new set of m ACs will be selected to run. Now, in order to maintain the comfort-band it should always be the case that the cumulative temperature rise of the $(n - m)$ ACs (OFF) is less than the cumulative temperature fall of the m ACs (ON) at $T_i = B^U$.

The rise and fall in temperatures of an ACs during a scheduling decision interval I^S can be expressed as $I^S \times S_r^i$ and $I^S \times S_f^i$ respectively. So, a sufficient condition for selecting any set of m ACs for the purpose of comfort-band maintenance is

$$m \times abs(S_f) \geq (n - m) \times S_r \tag{23}$$

where, the cooling slope $S_f = min(S_f^i \mid i = 1, 2 \cdots n)$ and the warming slope $S_r = max(S_r^i \mid i = 1, 2 \cdots n)$ at $T_i = B^U$. Note that if the Equation 23 is valid at $T_i = B^U$, it is also valid for any $T_i > B^U$, because of the exponential nature of the thermal characteristics of ACs.

Equation 23 is pessimistic, because temperature rising and falling slopes will be different for different ACs. The weaker sufficient condition is:

$$\sum_{i=1}^{m} abs(S_f^i) \geq \sum_{i=1}^{n-m} S_r^i \tag{24}$$

Where, i) the sum on the left is obtained from the first m of the n ACs arranged in ascending order of their falling slope and ii) the sum on the right is obtained from the first $n - m$ of the n ACs arranged in descending order according to their rising slope.

From Equation 10, we get

$$S_f^i \mid_{T_i = B^U} = \frac{dT_i}{dt} \mid_{T_i = B^U} = b_i(d_i - B^U) \tag{25}$$

where, b_i, d_i are constants corresponding to the i^{th} AC. Similarly from Equation 15, we get

$$S_r^i \mid_{T_i = B^U} = \frac{dT_i}{dt} \mid_{T_i = B^U} = b'_i(d'_i - B^U) \tag{26}$$

where, b'_i, d'_i are constants corresponding to the i^{th} AC. Substituting the values of S_f^i and S_r^i in Equation 24, the proof follows. ∎

IV. SIMULATION & PROTOTYPE IMPLEMENTATION OF TCBM

Here we report on the performance of our algorithm and compare the results with that of several candidate algorithms adapted from the real-time domain for AC scheduling. We also implemented TCBM algorithm in a prototype set-up and report on observations made from it.

Simulation studies were carried out based on thermal characteristics of ACs generated by curve-fitting using the empirical data presented in Section II-B. The load consisted of 5 ACs. We assume that peak power limit permits no more than 3 ACs to be ON at a time. Feasibility is checked according to Equation 22 for a comfort band with $T^U = 26^0C$ & $T^L = 23^0C$. Δ^U (0.24) and Δ^L (0.19) values are calculated according to the methods discussed in Section III-B and III-C respectively.

A. Candidate scheduling algorithms

We first applied global EDF (Earliest Deadline First) scheduling [10], Least Slack First (LSF) scheduling [2], and value-based scheduling [11] [12] from the real-time domain. The periodic task model as described in Section I, is used for AC tasks in this simulation with C_i as the time duration an AC needs to run to bring down T_i from T^U to T^L and L_i as the time duration it can be switched OFF, i.e., when T_i rises from T^L to T^U. The period P_i is its duty-cycle $(C_i + P_i)$.

In EDF [5], the task which has the earliest absolute deadline is scheduled first. Global EDF [10] extends the EDF scheduling of tasks to uniform multi-processors, where m out of n tasks can run on m processors at any point of time. Global EDF (gEDF) scheduling appears to be a natural candidate for AC scheduling as the constraint that only m out of n ACs can run simultaneously can be straight-away mapped to scheduling n tasks on m processors.

In LSF or *Least Slack-Time First (LST)* [6], the task which has the least slack (or laxity) is scheduled first, where at any time t, slack of a task having a deadline D_i is defined as $(D_i - t)$ minus the time required to complete the remaining portion of the task. In our case, slack at any time is the remaining length of time it can be OFF.

Value-based scheduling [11] [12] is another promising algorithm for scheduling TCE devices, as a TCED task can be associated with a value $V_i(T_i)$, depending on the existing temperature T_i of the environment under its control. Specifically, we can assign a value to the AC task according to the amount of cooling it can give per unit time when powered on and the desired comfort-band. The state-dependent attribute V_i has the highest value at $T_i = T^U$ and it becomes negative at $T_i > T^U$. V_i goes on decreasing as $(T^U - T_i)$ decreases and again attains negative value at $T_i < T^L$. We can summarize it as follows.

$$V_i(T_i) = \begin{cases} T^U - T_i, & \text{if } T_i > T^U , \\ \frac{T_i - T^L}{T^U - T^L}, & \text{if } T_i \leq T^U \end{cases} \quad (27)$$

B. Relative performance of the algorithms

Figures 3, 4, 5 and 6 show the simulation results of LSF, global EDF, Value-based and TCBM scheduling respectively, when ACs are started with a room temperature $T_r = 27.5^0C$. The observations are summarized in Table V. It is important to note that LSF, Value-based and TCBM scheduling take less than ≤ 10 minutes, which can be considered to be a reasonable amount of time, to bring down the room temperatures within $[26^0C, 23^0C]$ from the initial temperature of 27.5^0C and maintain it thereafter. In contrast, global EDF takes 31 minutes.

Global EDF does not always maintain the room temperature in the comfort-band, which is obviously unacceptable. Even under LSF scheduling, T_i goes beyond T^U on some occasions.

It can also be observed from Figure 3 that all ACs are switched OFF at different points of time causing higher T_i, because LSF policy keeps a load OFF irrespective of T_i, if it has consumed power (i.e., remains ON) for C_i amount of time within its period P_i. Similarly, in case of gEDF scheduling, in Figure 4, loads are switched OFF because they consumed power for C_i amount of time within their period P_i, without taking into consideration the value of T_i. This also explains the high values of discomfort duration shown in Table V under LSF and gEDF scheduling.

In case of TCBM, it can be observed in Figures 6 that m ACs are always ON, as long as they are required to maintain the comfort-bands. During the initial phase, when the zonal temperature is high (27.5^0C), T_i of some ACs increases only because it is not possible to run more than m ACs at a time.

The number of ON/OFF switchings of ACs in LSF and Value-based scheduling is very high, 400% and 1050% more respectively, compared to our TCBM scheduling. Switching in global EDF scheduling is 150% higher compared to TCBM scheduling. We calculated the expected number of switching of 5 ACs in 1500 min. of simulation considering 2 (ON & OFF) switchings per period when run only under thermostatic control. The number of excess switching against the minimum value of 492 is shown in Table V and it can be observed that the no. of excess switching under TCBM is only 70 as against 5507, 1705 & 384 under value-based, LSF and gEDF respectively.

In case of value based scheduling, it can be observed from Figure 5 and Table V that the number of AC switchings is extremely high (5999 for 5 ACs). It can be explained from the fact that as an AC is switched ON, its corresponding room temperature decreases, which in turn deceases the value of supplying power to it. In the subsequent scheduling period(s), the same AC is likely to be switched OFF because of its reduced value and hence later its value will rise again. Therefore, such cyclical rise and fall of values causes excessive switchings.

In Figure 7, we focus on a single AC. It highlights the superior performance of our TCBM algorithm with respect to different performance metrics.

C. Prototype Experimental Studies

The results of real-life control of 2 ACs in a room using TCBM algorithm, after both ACs reach the comfort-band, are shown in Figure 8. It may be noted here that, when an AC is switched ON, its compressor and fan run and when it reaches the lower value of its set point, the AC is switched to fan-mode (i.e., compressor is OFF, but the fan is running). [2]

Feasibility analysis according to Equation 24 indicates that comfort-band of $[26^0C - 23^0C]$ can be maintained with peak power permitting only one AC to be ON at a time. Figure

[2]The power requirement in the fan-mode is very low as compared to the same when AC is in cooling-mode (compressor & fan both running). Hence, for simplicity of explanation, we ignore the power consumed by the fan-mode.

	1500 Min.	$T^U = 26^0C$	$T^L = 23^0C$	
Metric	**LSF**	**gEDF**	**Value**	**TCBM**
Time to reach Comfort-band	10	Comfort-band not maintained	4	10
No. of Switching	2197	876	5999	562
Excess Switching	1705	384	5507	70
Discomfort duration	56	293	14	33

Fig. 3. LSF Scheduling of ACs

8 demonstrates that both ACs function alternatively, while maintaining the comfort-band. The values of Δ^U and Δ^L obtained using the methods discussed in Section III-B and III-C are 0.24^0C and 0.18^0C respectively.

In can be observed from Figure 8 that whenever an AC is switched ON, it can affect the temperature of the other zone marginally. For example, at time t=40, when AC1 is switched ON while the AC2 was OFF, the temperature of the zone corresponding to AC2 falls by about 0.2^0C and then goes up again. This is because our assumption that the zones controlled by the two ACs are thermally de-coupled may not hold true in reality, because they are in the same room. How to consider such dependencies is part of our ongoing work.

In summary, none of the candidate scheduling algorithms from the real-time domain namely, LSF, global EDF and Value-based scheduling, is suitable for TCED scheduling because they suffer from the following disadvantages:

- Maintaining room temperature within comfort zone $[T^U, T^L]$ is not guaranteed in LSF and gEDF scheduling, as the basic criterion for these algorithms is to provide resource for C_i units of time within every period P_i, irrespective of *when* this C_i time is allocated. This lacuna can cause the controlled environmental parameter to go beyond its limit even though the task gets resource (power in case of TCED) for C_i unit of time within every P_i, as explained in Section I.
- Excessive and undesirable switching of ACs since LSF, gEDF and Value-based policies decide to switch ON/OFF ACs irrespective of whether room T_i is within the comfort zone $[T^U, T^L]$.

Since it is important to avoid unnecessary switching (ON/OFF) as discussed in Section I, it is preferable that switching of an AC occurs only when its environmental parameter T_i is outside the specified comfort-band $[T^U, T^L]$. It may be argued that excessive switching under LSF, gEDF and Value-Based scheduling can be avoided if we simply apply those scheduling policies only when T_i is outside the comfort-band. But, switching (ON or OFF) of a device becomes mandatory if temperature goes beyond the comfort-band associated with the device. In other words, such a scheme leaves no scope for applying LSF, gEDF or Value-based scheduling once the T_i is outside the comfort band.
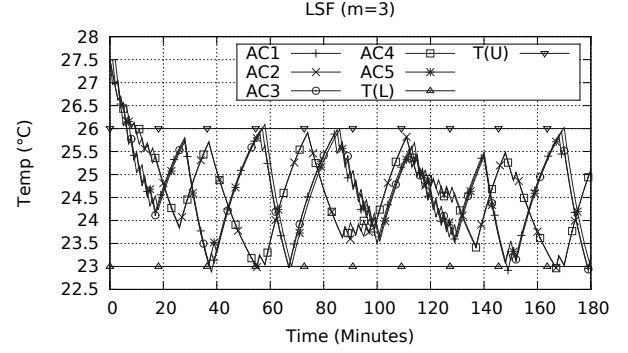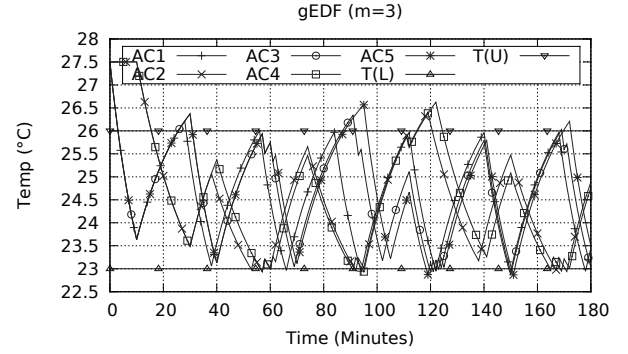

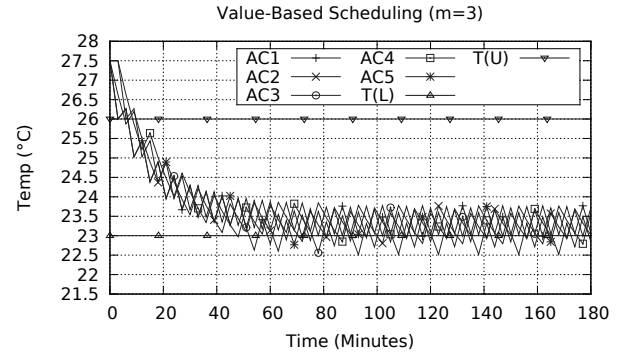Fig. 4. gEDF Scheduling of ACs

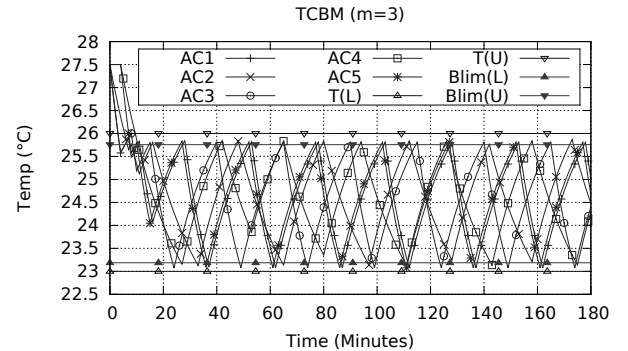Fig. 5. Value-Based Scheduling of ACs
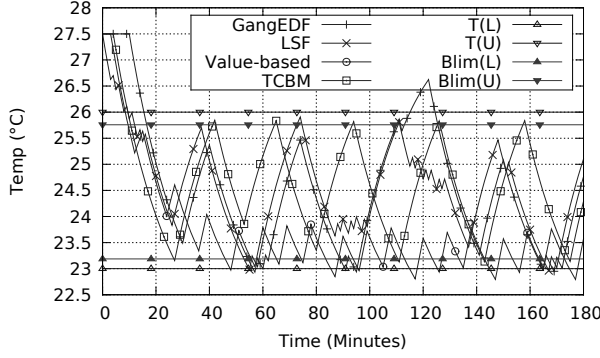

Fig. 6. TCBM Scheduling of ACs

Fig. 7. Effect of Different Scheduling Policies on an AC

Fig. 8. Effect of TCBM Scheduling on 2 ACs

V. ON-LINE ADAPTIVE CONTROL USING TCBM

We discussed in Section II-B that it may not be practical to calculate the constants in the AC thermal characteristic Equations 5 and 3 and suggested curve-fitting based on experimental data and obtain the constants (b, d & b', d') necessary for calculating Δ^U & Δ^L. We also discussed how change in comfort-band helps in reducing power consumption.

Therefore, we propose an adaptive technique for controlling TCE devices in order to address i) changes in the thermal characteristics of an AC under different ambient temperatures and ii) maintaining comfort-band under time varying peak power limit.

A. Handling varying ambient temperature

- Generate AC characteristics constants at various ambient temperature values (assuming change in heat loads to be negligible) obtained by off-line curve-fitting, store them in TCBM controller and use the same for adaptive control, by dynamically adjusting Δ^U & Δ^L.
- Measure ambient temperature and do on-line curve-fitting based on the room temperature data obtained in the immediate past. Generate values of Δ^U and Δ^L from the constants of the on-line curve-fit and apply TCBM algorithm. This scheme is capable of handling changes in heat loads.

TABLE VI
SCHEDULABILITY BY SHIFTING COMFORT-BAND (5 ACS)

Comfort-band (^0C)	m	$\sum S_f$	$\sum S_r$	feasibility ($\sum S_f \geq \sum S_r$)
$23-25$	3	0.72	0.34	Yes
$23-25$	2	0.35	0.53	No
$24-26$	2	0.46	0.35	Yes
$23-26$	2	0.47	0.35	Yes

B. Shifting/Relaxation of Comfort-Band

From Figure 2 (a) & (b), it can be observed that the duration for which AC remains ON is reduced if, i) the comfort-band is shifted up along the temperature scale or if ii) T^U is shifted up. This happens because of the fact that cooling rate of AC is more at higher temperatures. Now, from the feasibility condition (Equation 24), it can be observed that if the cooling slope S_f is increased (by shifting T^U up) feasibility condition will be satisfied with a lower value of m. Suppose, peak demand constraint allows power for 3 ACs at a time and we have 5 to control using TCBM algorithm. We consider that out of 5 ACs, two ACs have the same thermal characteristics as that of AC1 and three ACs have the same characteristics as that of AC2. The constants of their characteristic equations are shown in Table II and III. The data related to feasibility of these 5 ACs under TCBM are generated using Equations 25 & 26 are shown in Table VI.

It can be observed from Table VI, that it is feasible to maintain a comfort-band of $[23^0C - 25^0C]$ by running 3 ACs at a time. But, the same comfort-band can not be maintained, if peak demand constraint allows powering-ON of at most 2 ACs at a time.

Now, let's take two cases; i) comfort band is shifted from $[23^0C - 25^0C]$ to $[24^0C - 26^0C]$ and ii) upper limit of the comfort band is changed from 25^0C to 26^0C. It can be observed from Table VI that in both the cases, comfort-band can be maintained by running 2 out of 5 ACs at a time.

Therefore, we conclude that under varying peak demand limit, comfort-band $[T^U, T^L]$ can be suitably adjusted, as discussed above, to meet the peak demand constraint.

This has implications for adaptive demand-response in dynamic energy pricing/availability scenarios:

1) Given an externally-imposed peak demand constraint, the comfort-band can be adjusted dynamically and the user informed about it.
2) If a user insists on staying with a pre-set comfort-band, he can be warned about the implications of violating the peak power consumption limit ahead of time.

VI. RELATED WORK

Reduction in peak power demand for grid stability, reduction of occurrences of blackouts as well as equipment failures and improvement of energy efficiency have drawn the attention of many researchers. These cover the gamut from minimizing peak demand through buffering of renewable resources [13] to the use of model predictive control for energy efficiency

in building climate control [14]. The objective of our work is different since we focus on i) reduction of peak power demand by means of proper scheduling of background electrical loads keeping in mind practical considerations and ii) providing a feasibility criterion based on the requirement of thermal comfort-band maintenance.

Reduction in peak power demand using a global scheduler has been addressed in [15], [2] and [16]. Peak-power reduction by means of scheduling appliances using a combination of admission & curtailment control has been discussed in [15], but it does not discuss schedulability. A comprehensive study on power-usage in residential buildings has been discussed in [2] and Least Slack First (LSF) is proposed for background load scheduling. Lazy scheduling approach is proposed in [16] to control HVAC&R (heating, ventilation, air-conditioning & Refrigeration) devices for peak demand reduction, but it does not guarantee meeting the peak demand constraint. Modelling electrical loads, categorized under cyber-physical energy systems (CPES), as real-time tasks is discussed in [4] and [3] and application of traditional real-time scheduling algorithms like EDF are suggested. But, we have shown that they are not suitable for thermostatically controlled electrical devices (TCED), which may be considered as a special sub-class of CPES. Work on global EDF (gEDF) scheduling [10] is also related as scheduling of background loads can be directly mapped into scheduling of m out of n real-time tasks running in parallel on m uniform multiprocessors. Similarly, value-based scheduling discussed in [11] and [12] is related as value can be attached dynamically to a TCED task, based on the environmental parameter feedback. But, we showed that none of the LSF, gEDF and value-based policies are suitable to our problem as is, as either they involve unnecessary and undesirable switching of TCE devices, or do not guarantee maintenance of temperature within the comfort-band.

VII. Conclusion

In this paper, based on our empirical study of the functioning of ACs and by applying heat transfer principles, we developed a conceptual model of power consumption and thermal comfort-band maintenance. From the insights gained from this study, a feasibility analysis technique was proposed for maintaining thermal comfort-band under peak power demand constraint. We showed how this analysis can be utilized for adaptive demand-response control of TCE devices under time varying peak power demand constraint. Driven by the goal of maintaining the comfort-band for as many devices as possible with minimal number of switching of power between appliances, we presented the TCBM approach for selecting the subset of appliances to power at a given point in time.

Our performance study demonstrates the superior performance characteristics of our algorithm compared to algorithms adapted from the literature. It showed that existing scheduling algorithms for timely task execution are not suitable for scheduling TCE devices because, either i) they do not prevent undesirable switching (preemption) of the devices even when the temperature of the environment under their control is within the associated comfort-band or ii) they do not guarantee maintenance of comfort-band. In contrast, our TCBM algorithm offers a minimum of 400% less switching as compared to the candidate algorithms.

In this paper, ACs were chosen as a representative of the TCE devices. Generalizing this work, applying the lessons learned for other TCE devices and studying the thermal dependencies of zones controlled by multiple TCE devices are part of our ongoing work.

References

[1] TRF, *Understanding PECOs General Service Tariff.* Philadelphia, PA 19106-1591: (available at www.trfund.com), 2007.

[2] S. K. Barker, A. K. Mishra, D. E. Irwin, P. J. Shenoy, and J. R. Albrecht, "Smartcap: Flattening peak electricity demand in smart homes," in *Proceedings of the 2012 IEEE International Conference on Pervasive Computing and Communications*, March 2012, pp. 67–75.

[3] T. Facchinetti, E. Bini, and M. Bertogna, "Reducing the peak power through real-time scheduling techniques in cyber-physical energy systems," in *Proceedings of the 1st International Workshop on Energy Aware Design and Analysis of Cyber Physical Systems (WEA-CPS)*, April 2010.

[4] M. L. D. Vedova, E. D. Palma, and T. Facchinetti, "Electric loads as real-time tasks: An application of real-time physical systems," in *Proceedings of the 7th International Wireless Communications and Mobile Computing Conference, IWCMC 2011*, July 2011, pp. 1117–1123.

[5] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *J. ACM*, vol. 20, no. 1, pp. 46–61, 1973.

[6] J. W. S. Liu, *Real-Time Systems.* Pearson Education, 2000.

[7] P. C. Sen, *Principles of Electrical Machines and Power Electronics.* John Wiley and Sons (Asia) Pte. Ltd., Singapore, 2005.

[8] F. P. Incropera and D. P. DeWitt, *Introduction to Heat Transfer, Sixth Edition.* John Wiley and Sons, 2011.

[9] ASHRAE, *Fundamentals of HVAC Systems.* Atlanta, GA: ASHRAE Handbook, 2001.

[10] S. Baruah, "Techniques for multiprocessor global schedulability analysis," in *Proceedings of the 28th IEEE International Real-Time Systems Symposium, RTSS '07*, 2007, pp. 119–128.

[11] C. D. Locke, "Best-effort decision making for real-time scheduling," Ph.D. dissertation, Carnegie-Mellon University, Computer Science Department, 1986.

[12] A. Burns, D. Prasad, A. Bondavalli, F. D. Giandomenico, K. Ramamritham, J. A. Stankovic, and L. Strigini, "The meaning and role of value in scheduling flexible real-time systems," *Journal of Systems Architecture*, vol. 46, no. 4, pp. 305–325, 2000.

[13] A. Bar-Noy, M. P. Johnson, and O. Liu, "Peak shaving through resource buffering," in *Approximation and Online Algorithms,6th International Workshop, WAOA 2008*, September 2008, pp. 147–159.

[14] O. Frauke, P. b. Alessandra, C. N. Jonesc, G. Dimitrios, G. Markus, S. Vanessa, L. Beat, and M. Manfred, "Use of model predictive control and weather forecasts for energy efficient building climate control," *Energy and Buildings*, vol. 45, p. 1527, November 2012.

[15] G. Xiong, C. Chen, S. Kishore, and A. Yener, "Smart (in-home) power scheduling for demand response on the smart grid," in *IEEE Innovative Smart Grid Technologies (ISGT)*, January 2011, p. 17.

[16] T. Nghiem, M. Behl, R. Mangharam, and G. J. Pappas, "Green scheduling of control systems for peak demand reduction," in *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference, CDC-ECC 2011*, December 2011.