Automating occupant-building interaction via smart zoning of thermostatic loads
A switched self-tuning approach

Baldi, Simone; Korkas, Christos D.; Lyu, Maolong; Kosmatopoulos, Elias B.

Important note
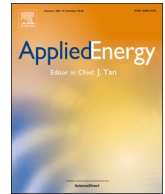To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Automating occupant-building interaction via smart zoning of thermostatic loads: A switched self-tuning approach

Simone Baldi[b,*], Christos D. Korkas[a,c], Maolong Lv[b], Elias B. Kosmatopoulos[a,c]

[a] *Dept. of Electrical and Computer Engineering, Democritus University of Thrace, Xanthi 67100, Greece*
[b] *Delft Center for Systems and Control, Delft University of Technology, Delft 2628CD, the Netherlands*
[c] *Informatics & Telematics Institute, Center for Research and Technology Hellas (ITI-CERTH), Thessaloniki 57001, Greece*

HIGHLIGHTS

- Development of a self-tuning load management program with smart zoning.
- Formulation of the control problem in a Hamilton-Jacobi-Bellman framework.
- Multiple behaviors triggered by set point rules are embedded in the optimization.
- Feedback strategies exploit information stemming from building and weather states.
- Simulations are conducted on an EnergyPlus model of an actual building in Greece.

ARTICLE INFO

ABSTRACT

Load management actions in large buildings are pre-programmed by field engineers/users in the form of if-then-else rules for the set point of the thermostat. This fixed set of actions prevents smart zoning, i.e. to dynamically regulate the set points in every room at different levels according to geometry, orientation and interaction among rooms caused by occupancy patterns. In this work we frame the problem of load management with smart zoning into a multiple-mode feedback-based optimal control problem: multiple-mode refers to embedding multiple behaviors (triggered by building-occupant dynamic interaction) into the optimization problem; feedback-based refers to adopting a Hamilton-Jacobi-Bellman framework, with closed-loop control strategies using information stemming from building and weather states. The framework is solved by parameterizing the candidate control strategies and by searching for the optimal strategy in an adaptive self-tuning way. To demonstrate the proposed approach, we employ an EnergyPlus model of an actual office building in Crete, Greece. Extensive tests show that the proposed solution is able to provide, dynamically and autonomously, dedicated set points levels in every room in such a way to optimize the whole building performance (exploitation of renewable energy sources with improved thermal comfort). As compared to pre-programmed (non-optimal) strategies, we show that smart zoning makes it is possible to save more than 15% energy consumption, with 25% increased thermal comfort. As compared to optimized strategies in which smart zoning is not implemented, smart zoning leads to additional 4% reduced energy and 8% improved comfort, demonstrating improved occupant-building interaction. Such improvements are motivated by the fact that the approach exploits the building dynamics as learned from feedback data. Moreover, the closed-loop feature of the approach makes it robust to variable weather conditions and occupancy schedules.

1. Introduction

The future will see more and more developments in the smart buildings and smart grids areas [1]: while smart buildings should implement *demand management* programs (sometimes also referred to as load management programs) [2], smart grids should implement *demand response* programs that can modify normal consumption patterns in buildings depending on the state of the grid [3]. It is well known that consumption patterns in buildings are deeply affected by heating, ventilating and air conditioning (HVAC) operation: almost 50% of the energy consumed by any building goes into HVAC [4]. In turn, HVAC operation is driven by the selection of thermostat set points. Therefore,

both load management programs and demand response programs should ultimately act on such set points. Two ways of acting on the set points are possible: in the first one, the users can actively manage the thermostat set points; in the second one, *automated* solutions for set point selection are established, in which the set point is automatically regulated without users being actively involved. However, automated demand management presents several challenges, one of the main ones being enhancing energy efficiency in thermostatically controlled HVAC loads via *smart zoning* [5]. Smart zoning is the capability of *dynamically* creating localized climate conditions that take into account the usage of a room, its orientation and its occupancy. For example, there is a recent trend in developing smart thermostats, Nest, Tado and Toon being just a few examples, that allow a sort of 'smart automated' regulation of the thermostat set point (e.g. via learning algorithms). While being appealing and in many case effective, such solutions work for small homes, and demand management actions often consist of a fixed (non-dynamic) set of rule-based options. This fixed set of options often neglects the building dynamics and the dynamic occupant-building interaction: in fact, in order to keep consistent performance, the HVAC set points should be continuously adjusted depending on variable weather conditions (which will affect the availability of renewable energy sources) or depending on user activity (which will affect occupancy patterns). A zoning program which is truly smart should (1) combine dynamically the available information stemming from the building and the weather states [6] (intelligent load management); (2) embed the occupancy pattern using models that can be clearly interpretable by human beings [7] (occupancy-based load management). This, unfortunately, turns out to be a big challenge due to the multiple factors that *dynamically* influence energy consumption [8]. The following subsections give an overview of recent results on intelligent load management and on occupancy-based load management. Some open problems in these areas are discussed, from which the motivations for this work arise.

1.1. Related work in intelligent load management

HVAC load management is the most cost-effective option for energy efficiency in buildings: while it is clear that raising the HVAC set point during summer and decreasing it during winter has great energy saving potential, smart zoning programs would push the energy efficiency even further by intelligently taking thermal comfort constraints into account [9]: the thermal conditions for human occupancy are codified in the ASHRAE Standard 55 [10]. Energy/comfort/economy costs are studied in [11] using a static building model. However, dynamics models are more appropriate to study the delicate trade-off between changing the thermostatic set point and the effect on thermal comfort: efforts in this direction can be found, for example in [12] via a simulation-based method, in [13] via weighted linguistic fuzzy rules in combination with a rule selection, in [14] via a numerical procedure based on the finite-difference method, and in [15] via population-based stochastic optimization, based on different comfort bounds. Being thermal comfort closely connected to indoor air quality, a conflict exists also between energy saving and indoor air quality improvement, as studied in [16] via a knowledge-based automation approach, or in [17] via a genetic algorithm. Sometimes the focus is on user satisfaction, which may or may not include thermal comfort: a rule-based demand side load management technique that is capable of controlling loads within the residential building in such a way that the user satisfaction is maximized is considered in [18].

Importance of a dynamic optimization: With a few exception, most state-of-the-art works consider the optimization of 'static' parameters: such parameters are not able to evolve dynamically and in real-time if new conditions arise. In order to achieve dynamic optimization, feedback-based strategies are necessary. This is particularly relevant if renewable energy sources like PV panels must be exploited [19], so that the HVAC management should take this information into account to minimize non-renewable energy consumption. In fact, a reasonable criterion in addition to thermal comfort/user satisfaction, is the one of covering the HVAC load using renewable energy sources: [20] tackles the problem via a rule-based algorithm that controls the battery inverter, whereas [21] considers different rule-based strategy planning models that allow to select optimum preheating/cooling time. When considering the power demand of aggregated HVAC, the HVAC control problem is formulated as a scheduling problem in [22]. Few works consider dynamic programs for joint energy savings with thermal comfort, and no works, to the best of the authors' knowledge, embed smart zoning in such programs. For example, in previous work by some of the authors, a combined criterion composed of the non-renewable energy consumption and the thermal comfort has been used in [23] to design an appropriate feedback-based strategy: however, the fact that the availability of the building testbed was limited to a few rooms prevented the implementation of a true zoning strategy.

1.2. Related work in occupancy-based load management

No work of the one cited considers occupancy-based strategies, i.e. setting the temperature in every room depending on occupancy patterns. Among the few works available in literature, [24] has considered selecting different set points based on aligning the residents' thermostat preferences with the indoor temperature, whereas [25] has considered aligning building resident's thermal preferences by assigning optimal resident-apartment pairs via integer-programming. The system proposed in [26] aims to match thermal service with the spatial distribution of occupants. In general, the goal of these works is to minimize the difference between unregulated room/zone temperature and the occupants' thermal preference based on heating/cooling loads: therefore, no dynamic smart zoning depending on weather conditions or occupancy schedule is considered.

Importance of a dynamic optimization: The importance of dynamically exploiting occupancy information in open-loop solutions like model predictive control has been recognized as a key enabler for energy efficiency with thermal comfort [27]. Interestingly, in [28] an occupancy-based rule-based controller is compared with an occupancy-based model predictive control that requires real-time optimization: it is found that the much higher complexity of the model predictive control yields negligible benefits over the simple rule-based controller. This suggests that feedback-based (closed-loop) solutions are of fundamental importance for occupancy-based load management. The integration of (rule-based) expert knowledge with automated feedback-based decisions is not trivial, since most advanced programs based on model predictive control can achieve this either by resorting to complex mixed-integer nonlinear programming [29] or by adding new constraints [30], which might make the optimization infeasible. Previous work by some of the authors has shown that occupancy information can be embedded not only in open-loop solutions, but also in closed-loop solutions [31]: furthermore, it was shown that occupancy information can be efficiently combined with availability of renewable energy supply so as to shape the demand based on the real-time building/weather/occupants measurements [32]. However, in such works the occupancy schedule is homogeneous within a building: no works, to the best of the authors' knowledge, explore the possibility of embedding the occupant behavior to implement zoning strategies that can take actions in a systematic and dynamic way based on the real-time building/weather/occupants measurements.

1.3. Motivations and contributions of this work

Summarizing, the motivations for us to implement an automated smart zoning program can be listed as:

- To truly optimize the energy consumption and thermal comfort, it is crucial to consider the 'system-of-systems' structure of buildings

(composition of interacting rooms, with dense interconnection of HVAC actuation/sensing [33]);

- When delivering their load management actions, smart thermostats still cannot automatically account for user behavior and occupancy patterns at the zone level (with different actions in different zones [34]);
- Rule-based actions alone cannot not promote dynamic set point adjustment based on internal conditions (building state) or external conditions (weather state) [35].

We tackle the aforementioned difficulties by embedding the load management with smart zoning program into a Hamilton-Jacobi-Bellman (HJB) optimal control problem whose main components can be identified as:

(a) Closed-loop control: by using feedback information stemming from the internal building state (temperatures, occupancy schedule, availability of renewable energy) and from the external state (weather conditions, weather forecasts) we are able to intelligently adapt the HVAC set point to all these conditions;
(b) Multi-modal control: by integrating the user behavior via a switched model that includes the occupancy state of the different rooms, we can explicitly consider the occupant-building interaction and the user-driven energy pattern;
(c) Self-tuning control: by solving the HJB framework via parameterization the candidate solution, we are able to use learning mechanisms to search for the optimal solution in an adaptive self-tuning way.

We refer to the proposed approach as R-PCAO (Rule-based Parameterized Cognitive Optimization).

The paper is organized as follows: Section 2 introduces the problem setting, while Section 3 focuses on the control goals. Section 4 presents the automated load management, with simulations performed in Section 5 to demonstrate and analyze performance. Conclusions are in Section 6.

2. Problem setting

The main purpose of HVAC automated programs should be to employ feedback to establish meaningful 'relations' among data gathered from internal states (temperature in each room, occupancy schedule) and from external states (weather conditions). Such relations should be ultimately exploited to select the HVAC set point in each room. In this section we present the mathematical models describing the 'relations' among different states. We consider a cooling problem and we start from the dynamics of a single room i with thermostatically controlled HVAC

$$
C_R \dot{T}^{(i)} = \begin{cases} \alpha(T_O - T^{(i)}) + \beta(T^{(\underline{i})} - T^{(i)}) \\ \quad + \gamma S_O + Q_{occ}^{(i)} + Q_{HVAC}^{(i)}, & \text{if } T^{(i)} \geqslant T_{set}^{(i)} \\ \alpha(T_O - T^{(i)}) + \beta(T^{(\underline{i})} - T^{(i)}) \\ \quad + \gamma S_O + Q_{occ}^{(i)}, & \text{if } T^{(i)} < T_{set}^{(i)} \end{cases} \tag{1}
$$

In (1), the superscript (i) is used to indicate a quantity of room i; the superscript $(\underline{i})$ is used to indicates a quantity of neighboring room[1] of room i. The other parameters are as follows: C_R is the thermal capacitance of the room, $T^{(i)}$ the room temperature, $T^{(\underline{i})}$ the neighbor room temperature, T_O the outside temperature, S_O the solar gain entering from the window. The parameters α and β are material-dependent and room-dependent: in particular, they indicate the heat transfer

coefficients between the room and the outside, and between the room and the neighboring rooms, respectively. The parameter γ depends on the size of the window (the larger the window, the larger the solar radiation entering the room), and might also be room-dependent.[2] Two other inputs affect (1): the first is the heat gain $Q_{occ}^{(i)}$ resulting from the presence of occupants (which clearly cannot be controlled), and second is the power $Q_{HVAC}^{(i)}$ injected by the HVAC to cool the room (which can be controlled). The mechanism for controlling the HVAC is the classical thermostatic mechanism, whose switching is driven by $T^{(i)}$ and $T_{set}^{(i)}$ as illustrated in (1).

Note that the effect of thermostatic threshold can be removed in the presence of a variable-speed drive. In fact, variable-speed drives can modulate the HVAC action in such a way that: the more the difference between the set point and the room temperature, the more the power injected by the HVAC; the less the difference between the set point and the room temperature, the less the power injected by the HVAC. To model such variable-speed drive, we consider

$$
Q_{HVAC}^{(i)} = K(T_{set}^{(i)} - T^{(i)}) \tag{2}
$$

where K is the constant of the proportional controller in the variable-speed drive. Till now we have considered a generic set point $T_{set}^{(i)}$: however, in practice the set point is scheduled by the building management system according to some rule-based strategies: in the following we describe the strategy according to which the set point is scheduled.

2.1. Rule-based set point selection

It is common practice of many office buildings to deploy a rule-based load management. Basically, the rules determine the HVAC set points based on the occupancy schedule. In this work we focus on three basic rules, implemented nowadays in the majority of the building/facility management systems:

1. **Normal mode**: This represents the desired set point when there are people in the room.
2. **Set-back mode**: This is the most common strategy in buildings. The load management program selects a higher set point during non-occupancy hours (e.g. outside office hours). In fact, it is usually preferable to switch-off HVACs when occupants are not present, aiming for lower energy consumption.
3. **Pre-cooling mode**: A common feature of many load management programs is to turn on the HVAC some time before people arrival, usually with a set point slightly lower than normal (for cooling) in order to reach faster appropriate indoor conditions.

Next to these three rules, a so-called **zoning program** can allow to select a different set point for every room. This might be necessary, e.g. due to different usage of a room, different window area and orientation, or even different preferences of the users. Note that in large buildings like office buildings or commercial buildings the difference between the set points in the different rooms can be 2–3 °C, therefore very relevant. We assume that such a program is available in our test case in view of better energy efficiency.

It is clear that the switch from one mode to another is driven by the occupant behavior, as represented in Fig. 1. Of course, the implementation of the schedule in Fig. 1 requires the presence of a system with the ability of predicting user behavior: this is a topic of increasing interest in recent years, cf. the survey [36] and the work [37]. Therefore, we assume the presence of such a prediction system.

---

[1] In case a room has more than one neighboring rooms, one should consider the summation of all these terms. In order to avoid making the notation more cumbersome, in the following we will not report such summation for simplicity.

[2] Therefore, α, β and γ should actually be $\alpha^{(i)}$, $\beta^{(i)}$ and $\gamma^{(i)}$, even of this is not explicitly indicated for compactness.

$$C_R \dot{T}^{(i)} = \begin{cases} (-\alpha-\beta + K)T^{(i)} + \beta T^{(\dot{\iota})} & \text{if } T^{(i)} \geqslant T^{(i)}_{set\text{-}point} \quad \text{and MODE} \\ \quad + \alpha T_0 + \gamma S_O + Q^{(i)}_{occ} & \qquad\qquad = \text{normal} \\ \quad -K T^{(i)}_{set\text{-}point}, \\ (-\alpha-\beta + K)T^{(i)} + \beta T^{(\dot{\iota})} & \text{if } T^{(i)} < T^{(i)}_{set\text{-}point} \quad \text{and MODE} \\ \quad + \alpha T_0 + \gamma S_O + Q^{(i)}_{occ}, & \qquad\qquad = \text{normal} \\ (-\alpha-\beta + K)T^{(i)} + \beta T^{(\dot{\iota})} & \text{if } T^{(i)} \geqslant T^{(i)}_{set\text{-}back} \quad \text{and MODE} \\ \quad + \alpha T_0 + \gamma S_O & \qquad\qquad = \text{set-back} \\ \quad -K T^{(i)}_{set\text{-}back}, \\ (-\alpha-\beta + K)T^{(i)} + \beta T^{(\dot{\iota})} & \text{if } T^{(i)} < T^{(i)}_{set\text{-}back} \quad \text{and MODE} \\ \quad + \alpha T_0 + \gamma S_O, & \qquad\qquad = \text{set-back} \\ (-\alpha-\beta + K)T^{(i)} + \beta T^{(\dot{\iota})} & \text{if } T^{(i)} \geqslant T_{pre\text{-}cooling} \quad \text{and MODE} \\ \quad + \alpha T_0 + \gamma S_O & \qquad\qquad = \text{pre-cooling} \\ \quad -K T^{(i)}_{pre\text{-}cooling}, \\ (-\alpha-\beta + K)T^{(i)} + \beta T^{(\dot{\iota})} & \text{if } T^{(i)} < T^{(i)}_{pre\text{-}cooling} \quad \text{and MODE} \\ \quad + \alpha T_0 + \gamma S_O, & \qquad\qquad = \text{pre-cooling} \end{cases}$$

(3)

$$\pi^{(i)} = \begin{cases} K(T_{set\text{-}point}-T^{(i)}) & \text{if } T^{(i)} \geqslant T_{set\text{-}point} \quad \text{and MODE = normal} \\ \quad -k S_O \\ \quad + \rho \\ \quad (T_{set\text{-}point}-T^{(i)})^2, \\ -k S_O & \text{if } T^{(i)} < T_{set\text{-}point} \quad \text{and MODE = normal} \\ \quad + \rho \\ \quad (T_{set\text{-}point}-T^{(i)})^2, \\ K(T_{set\text{-}back}-T^{(i)}) & \text{if } T^{(i)} \geqslant T_{set\text{-}back} \quad \text{and MODE = set-back} \\ \quad -k S_O \\ \quad + \rho \\ \quad (T_{set\text{-}back}-T^{(i)})^2, \\ -k S_O & \text{if } T^{(i)} < T_{set\text{-}back} \quad \text{and MODE = set-back} \\ \quad + \rho \\ \quad (T_{set\text{-}back}-T^{(i)})^2, \\ K(T_{pre\text{-}cooling}-T^{(i)}) & \text{if } T^{(i)} \geqslant T_{pre\text{-}cooling} \quad \text{and MODE} \\ \quad -k S_O & \qquad\qquad = \text{pre-cooling} \\ \quad + \rho \\ \quad (T_{pre\text{-}cooling}-T^{(i)})^2, \\ -k S_O & \text{if } T^{(i)} < T_{pre\text{-}cooling} \quad \text{and MODE} \\ \quad + \rho & \qquad\qquad = \text{pre-cooling} \\ \quad (T_{pre\text{-}cooling}-T^{(i)})^2, \end{cases}$$
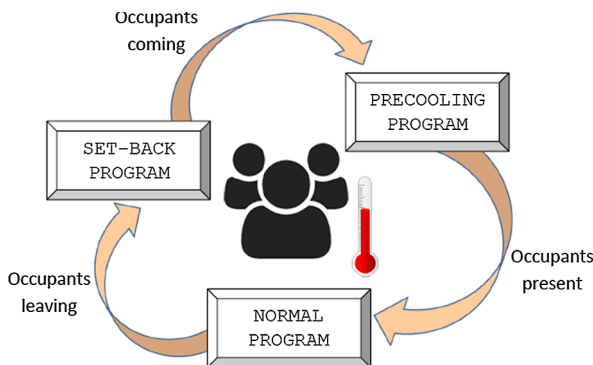
(4)

Fig. 1. Dynamics for set point selection in a single room.

$$\dot{T}^{(i)}(t) = A^{(i)}_{\sigma^{(i)}(t)} \underbrace{\begin{bmatrix} T^{(i)}(t) \\ T^{(\dot{\iota})}(t) \end{bmatrix}}_{T^{(i)}(t) = x^{(i)}(t)} + B^{(i)}_{\sigma^{(i)}(t)} \underbrace{T^{(i)}_{set\,\sigma^{(i)}(t)}(t)}_{u^{(i)}_{\sigma^{(i)}(t)}(t)}$$
$$+ E^{(i)}_{\sigma^{(i)}(t)} \underbrace{\begin{bmatrix} T_O(t) \\ S_O(t) \end{bmatrix}}_{d(t)},$$
$$\pi^{(i)}_{\sigma^{(i)}(t)}(t) = \pi^{(i)}\left(T^{(i)}(t), T^{(i)}_{set\,\sigma^{(i)}(t)}(t)\right)$$

(5)

2.2. Overall simplified building model

By combining the single room dynamics with the set point selection, one obtains a dynamical model represented in Fig. 2. This model is represented in analytical form in (3)–(5). The following points should be highlighted in the model:

- Let N_R be the number of rooms. As explained, each room i can operate in three modes, described by the index $\sigma^{(i)}(t)$: normal mode, set-back mode, and pre-cooling mode.
- Each room i has one state (its own temperature $T^{(i)}$) and one input (its own set point, $T^{(i)}_{set\text{-}point}$, $T^{(i)}_{set\text{-}back}$ or $T^{(i)}_{pre\text{-}cooling}$, depending on the active mode).
- Four external disturbances act on each room:
  - The term $T^{(\dot{\iota})}$ from a neighboring room of room i: this term acts as a disturbance for room i, but it is another one of the states of the building;
  - The outside temperature T_O;
  - An solar radiation S_O;
  - The occupancy heat gain $Q^{(i)}_{occ}$.
- It is assumed that forecasts for both T_O and S_O are available, so that these disturbances will be treated as present and future measurable disturbances.
- The occupancy schedule leading to $Q^{(i)}_{occ}$ is also assumed to be known at each time.
- The cost π comprises two terms:
  - The power provided by the HVAC, $K(T^{(i)}_{set}-T^{(i)})$ (active only when the HVAC is ON) minus the solar energy $k S_O$ collected by the solar panel (where k is the efficiency and size of the solar panel);
  - A weighted deviation from the desired set point which is used to model thermal comfort: $\rho(T^{(i)}_{set}-T^{(i)})^2$, where $T^{(i)}_{set}$ depends on the active mode.
- By taking into account the building topology, the model can be extended from one room to the entire building, as shown in (5). We do not describe the complete procedure for lack of space: the attentive reader will recognize that the following model will be obtained

$$\dot{x}(t) = A_{\sigma(t)}x(t) + B_{\sigma(t)}u(t) + E_{\sigma(t)}d(t)$$

(6)

$$\Pi_{\sigma(t)}\left(x(t), u(t)\right) = \sum_{i=1}^{N_R} \pi^{(i)}_{\sigma(t)}\left(x^{(i)}(t), u^{(i)}(t)\right)$$

(7)

where $x = [T^{(1)}, ..., T^{(N_R)}]'$, $u = [T^{(1)}_{set}, ..., T^{(N_R)}_{set}]'$, $d = [T_O \ S_O]'$ and $\sigma = col(\sigma^{(1)}, ..., \sigma^{(N_R)})$. Note that each room has its own switching signal because its is independent from the modes of the other rooms. In addition, the matrices $A_{\sigma(t)}, B_{\sigma(t)}, E_{\sigma(t)}$ arise from the connection
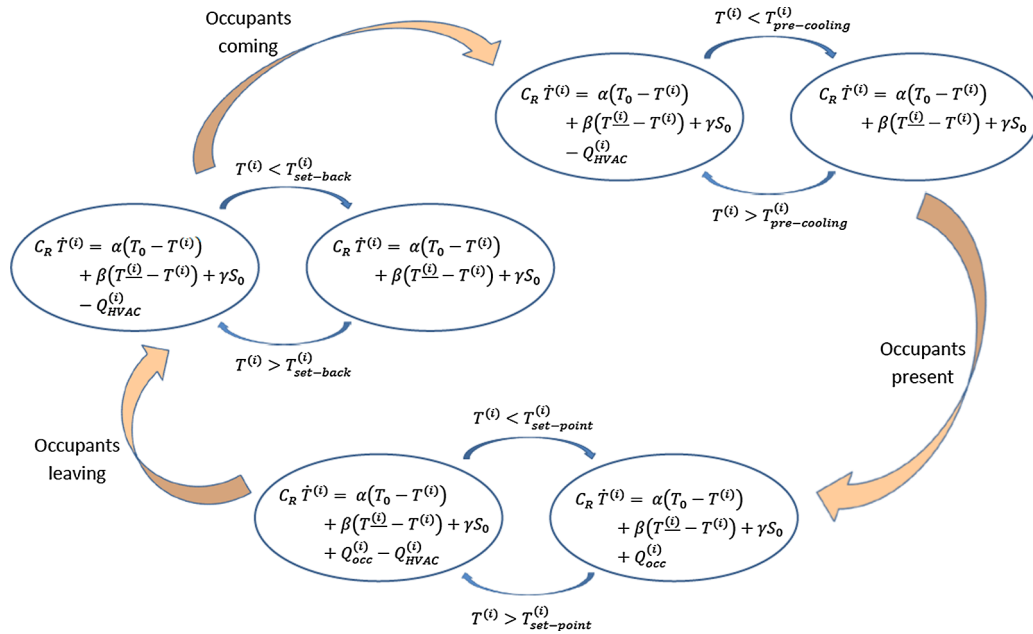
Fig. 2. Switched dynamics for one room (room i) with set point selection.

of all rooms and Π is the summation of the costs in each room.

3. Control goals

The specific goal of any load management program is to reduce the energy costs while keeping occupants satisfied, as formalized in (5). However, (5) refers to a specific instant in time, whereas, due to the *dynamic* evolution of temperatures, solar radiation and occupancy schedules (at least at the time scale of minutes), one should integrate the cost (5) over some long enough horizon

$$\int_0^{N_t} \Pi_{\sigma(t)}(t)dt = \int_0^{N_t} \sum_{i=1}^{N_R} \pi_{\sigma(t)}^{(i)}\left(T^{(i)}(t), T_{set_{\sigma(t)}}^{(i)}(t)\right)dt \tag{8}$$

where N_t is the horizon length, which can be anything from one day or one week, or even longer depending on the length of the simulations one is interested in performing. In many cases, people is interested in infinitely long horizons, with a discounting factor so as to obtain a finite integral

$$J = \int_0^{\infty} e^{-\lambda t}\Pi_{\sigma(t)}(t)dt \tag{9}$$

where $\lambda > 0$ is the discounting factor that reduces the importance of the cost far in the future. Clearly, the importance of minimizing an integral cost (9) in place of an instantaneous one like (5), is that the minimization of an integral cost takes into account dynamic behaviors in conjunction with occupancy schedule and weather conditions (an example is to 'overcool' the building when enough solar energy is available, so that energy can be saved when less solar energy is available [19]).

Summarizing, a dynamic optimization problem is obtained

$$\min J = \int_0^{\infty} e^{-\lambda t}\Pi_{\sigma(t)}(t)dt \tag{10}$$

$$s.t.$$
$$\dot{x}(t) = A_{\sigma(t)}x(t) + B_{\sigma(t)}u_{\sigma(t)}(t) + E_{\sigma(t)}d(t) \tag{11}$$

where all the variables have been defined in (3)–(7). Let us underline that the dynamics in (11) describe dynamics in the scale of minutes, which is quite crucial in order to account for changes in weather conditions and changes in occupancy. Clearly, the model (11) is a simplified room/building model: it is useful to define the state and inputs of

the system and the control objective, but it cannot be used for realistic testing of a smart zoning program: it is well known that simulation tools like EnergyPlus or TRNSYS [38] provide more realistic building dynamics. Nevertheless, the simplified dynamics summarized in (11) are of fundamental importance to understand which measurements from EnergyPlus or TRNSYS can actually be used by the optimization algorithm for feedback and real-time control.

3.1. Simulation model

To test the proposed algorithm on a realistic model, we use a building test case in EnergyPlus [39]. Our EnergyPlus model, shown later in Fig. 3, represents an actual office building in the campus of Technical University of Crete, Greece. The building has 10 rooms with 10 different HVAC set points that can be selected independently. Interestingly, the building is oriented along the North-South axis, with offices on each side of the building: note that, in view of its orientation, the offices take considerably different solar radiation from their windows. As a result, the different solar gain might influence drastically the selection of the HVAC set point. The buildings is also equipped with a photovoltaic panel that can be used to partially cover the energy demand. The EnergyPlus model has been developed during previous European research projects, mainly the AGILE project [40], coordinated by one of the authors, prof. Kosmatopoulos. It has been developed and validated in such a way that the thermal and energy dynamics of the model can capture in a realistic way the actual dynamics of the building. The energy cost and the comfort cost are automatically calculated by EnergyPlus. Because the photovoltaic energy is free of charge, the total energy cost (in kWh) takes into account only the energy absorbed from the power grid: in other words, if not further explained, in the following we will use the term 'energy consumption' to indicate the non-renewable portion of the energy consumption. To make the simulation even more realistic, typical load profiles from the actual buildings (PCs and appliances) have been implemented in EnergyPlus: clearly these loads are uncontrollable, but they make the total energy consumption of the building more realistic.

For the thermal comfort cost, we resort to an established metric, standardized in the ANSI/ASHRAE Standard 55 [10]: the Predicted Mean Vote Index (PMV) index. The PMV index is a thermal comfort model predicting the mean response of people according a seven-grade
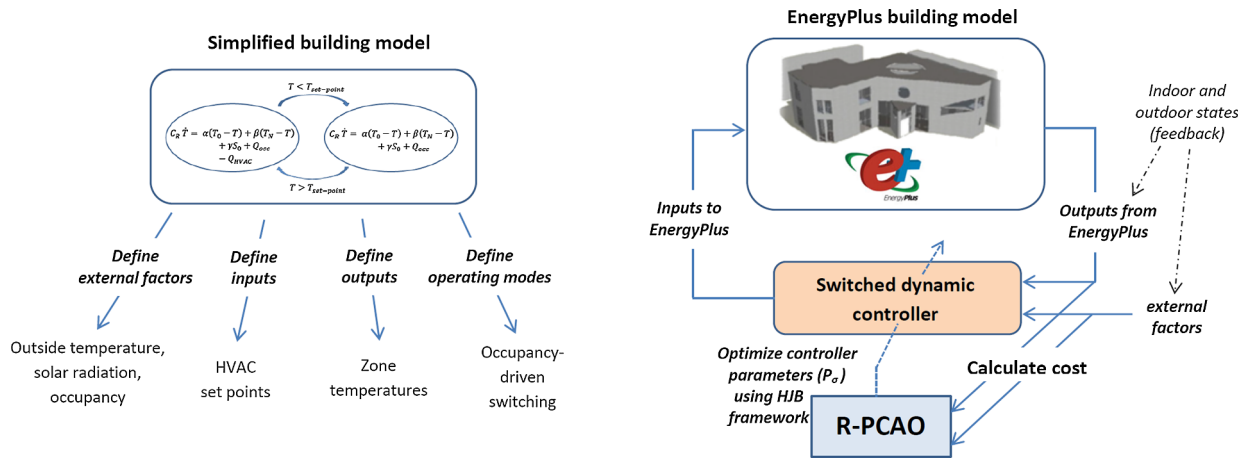
Fig. 3. The connection between the simplified building model, the EnergyPlus model and the optimization algorithm.

thermal sensation scale (from hot to cold). The PMV mean response is transformed via the Fanger's equation the Predicted Percentage of Dissatisfied people (PPD), expressed in %. The Fanger's equation is numerically solved by EnergyPlus.

3.2. Comparison strategies

For comparison purposes, the following load management strategies are adopted and implemented in EnergyPlus:

- Two Fixed Set Point (FSP) strategies. The FSPs employ a simple strategy, which consists of fixing the HVAC set points of each room at 24 °C ($FSP_{24°C}$) or 25 °C ($FSP_{25°C}$) during occupancy hours (the set points are 30 °C outside occupancy hours, which implies switching off the HVAC). Such simple strategies (they actually implement no zoning) provide acceptable performances in terms of the cost (10), although the performance is clearly far from optimal: the choice of these two set point temperatures is motivated by the AGILE project. During the project, fixed set point temperatures have been studied in such a way to find a trade-off between good energy consumption and good thermal comfort. It turned out that, in summer season, a set point of 24 °C gives an acceptable PPD of 10–12% (the ASHRAE standard suggests a PPD of around 10%), whereas a set point of 25 °C makes the PPD 1–2% worse while reducing the energy consumption of around 20%. Thus, these two strategies can be considered as two extremes of the Pareto front in between which optimal control strategies can play: most importantly, such strategies also provide with a fair base scenario that reduce any bias arising from calculating improvements for different weather conditions.
- A Rule Based Load Management (RBLM), with the following set of rules: set points constant to $T_s^{(i)}$ °C in room i during normal mode; set points of each room constant to 30 °C during set-back mode (which implies switching off the HVAC); set points constant to $T_s^{(i)}-1$ °C in room i half an hour before people entering the room (i.e. pre-cooling mode). Such pre-cooling strategy is also motivated by AGILE project [40], as a trade-off between the two FSP strategies. The values $T_s^{(i)}$°C are optimized with a genetic algorithm in such a way that the PPD is below 9% for a nominal occupancy schedule and some nominal weather conditions. This guarantees to have a baseline strategy that keeps occupants satisfied, so that we can perform meaningful comparisons with regards to energy consumption.

It is important to remark that the good performance for RBLM can only be evaluated under nominal conditions (weather and occupancy): this is because the designed RBLM is an open-loop schedule whose optimization would have to run continuously, otherwise its performance cannot be robust to changing conditions (e.g. weather and occupancy conditions). From these considerations we infer that a truly optimal strategy should be feedback-based, able to exploit information stemming from the entire building, and, in particular, able to use the information about environmental conditions so as to continuously and automatically adjust the set points: our proposed solution follows this idea as explained in the next section.

4. The proposed optimization methodology

Here we present the Rule-based Parameterized Cognitive Adaptive Optimization (R-PCAO) we adopted to solve the load management with smart zoning problem defined by (10) and (11). We will first give a brief overview of the concept, and then we will provide the details of the optimization algorithm.

4.1. The concept

The connection between the simplified building model, the EnergyPlus model and the optimization algorithm are shown in Fig. 3. Basically, the components interact as follows:

- The *simplified building model*: defined by (10) and (11), it is used to identify the inputs and outputs of interest (HVAC set points and zone temperatures, respectively), as well as the external factors influencing temperature (outside temperature, solar radiation, occupancy) and the different operating point of the HVAC (trigger by the occupancy pattern).
- The *EnergyPlus building model*: it provides the actual measurements to be used for feedback. To perform the simulation tests the EnergyPlus building model receives the inputs from the controller.
- The *switched dynamic controller*: it processes all the quantities previously identified (building and weather states) in such a way to generate the feedback-based input actions. The controller depends on some parameters (the control gains to be introduced later).
- The *self-tuning optimizer*: it must be able to self-tune the controller parameters such a way to maximize the building performance. This is the task of the proposed Rule-based Parameterized Cognitive Adaptive Optimization (R-PCAO) which exploits an Hamilton-Jacobi-Bellman (HJB) formulation described in the next section.

Depending on the building structure (defined by the simplified Eqs. (10) and (11), or even by the more realistic EnergyPlus model), different control gains are needed to optimize the performance at the whole building level. Therefore, one requires an optimization algorithm with the ability to 'learn' the building dynamics from data, and use this

Table 1
Quantities involved in the R-PCAO methodology.

Explanation	Symbol
State	x
Input	u
Mode	σ
State matrix	A_σ
Input matrix	B_σ
Exogenous matrix	E_σ
Instantaneous total cost	Π_σ
Instantaneous state cost	H_σ
Instantaneous input cost	$\bar{\rho}$
Discounting factor	λ
Optimal value function	$V^\circ{}_\sigma$
Optimal input	$u^\circ{}_\sigma$
Estimated value function	$\widehat{V}_\sigma$
Estimated optimal input	$\widehat{u}_\sigma$
Transformed state	z
Jacobian of z wrt x	M_z
Optimal parameterization	$P^\circ{}_\sigma$
Estimated optimal parameterization	$\widehat{P}_\sigma$
Candidate optimal parameterization	$\widehat{P}_\sigma^{(cand)}$
Best candidate optimal parameterization	$\widehat{P}_\sigma^{(best)}$
Candidate perturbation	$\Delta \widehat{P}_\sigma^{(i)}$
Approximation error	$O(1/L)$
Close-to-optimality index	ϵ
Estimated close-to-optimality index	$\widehat{\epsilon}$
Gradient update gain	a

knowledge to optimize the whole building performance (in our case, exploit the renewable energy sources while delivering improved thermal comfort). Before giving the mathematical details of R-PCAO, it is instructive to collect in Table 1 all the quantities involved in the optimization.

4.2. The Rule-based Parameterized Cognitive Adaptive Optimization

For convenience, let us first separate the cost $\Pi_{\sigma(t)}(t)$ in (10) as

$$\int_0^\infty e^{-\lambda t} \Pi_{\sigma(t)}(t) dt =$$
$$\int_0^\infty e^{-\lambda t} [H_{\sigma(t)}(x(t), d(t)) + \bar{\rho} u_{\sigma(t)}{}'(t) u_{\sigma(t)}(t)] dt \qquad (12)$$

for some appropriate $\bar{\rho} > 0$. From dynamic programming theory [41], we know that the optimal solution to (10) and (11) satisfies the Hamilton-Jacobi-Bellman equation:

$$u_{\sigma(t)}{}^\circ(x(t)) =$$
$$\underset{u}{\mathrm{argmin}} \left\{ \left(\frac{\partial V^\circ{}_{\sigma(t)}(x)}{\partial x} \right)' (A_{\sigma(t)} x(t) + B_{\sigma(t)} u_{\sigma(t)}(t) + E_{\sigma(t)} d(t)) \right.$$
$$\left. + H_{\sigma(t)}(x(t), d(t)) + \bar{\rho} u_{\sigma(t)}{}'(t) u_{\sigma(t)}(t) - \lambda V^\circ{}_{\sigma(t)}(x(t)) \right\} \qquad (13)$$

where $V^\circ{}_\sigma$ is typically referred to as the optimal value function, while $u^\circ{}_\sigma$ is referred to as the optimal control. Note that the optimal value function and the optimal control are dependent of $\sigma(t)$, in view of the different modes of the system. The main idea behind the R-PCAO algorithm is to parameterize both the optimal value function and the optimal control

$$V^\circ{}_{\sigma(t)} = z(x(t))' P^\circ{}_{\sigma(t)} z(x(t)) + O(1/L) \qquad (14)$$

$$u^\circ{}_{\sigma(t)} = -\bar{\rho}^{-1} B_{\sigma(t)}{}' M_z(x) P^\circ{}_{\sigma(t)} z(x(t)) + O(1/L) \qquad (15)$$

where $P^\circ{}_\sigma$ can be referred to as the optimal (or nearly optimal) parameterization matrix. Because the value function can be interpreted as a Lyapunov function for the system, it is positive definite, which can be achieved by imposing $\epsilon_1 I \leqslant P^\circ{}_\sigma \leqslant \epsilon_2 I$. In addition, the function $M_z(x)$ is the Jacobian matrix of $z(x)$ with respect to x, $z(x)$ is the feedback vector to be defined later, and $O(1/L)$ is the approximation term. The exact form of L and $z(x)$ will be discussed later: for the moment it is sufficient to say that L is a parameter such that by increasing L the approximation error becomes smaller (similar to neural-network approximation error). The form of $z(x)$ also depends on L (similar to neural-network regressors).

Being $V^\circ{}_{\sigma(t)}$ unknown, the main problem resides in the fact that the optimal parameterization $P^\circ{}_\sigma$ is unknown. However, one can substitute the optimal parameterization matrix $P^\circ{}_\sigma$ with an estimate $\widehat{P}_\sigma$, with $\epsilon_1 I \leqslant \widehat{P}_\sigma(t) \leqslant \epsilon_2 I$

$$\widehat{V}_{\sigma(t)} = z(x(t))' \widehat{P}_{\sigma(t)}(t) z(x(t)) \qquad (16)$$

$$\widehat{u}_{\sigma(t)} = -\bar{\rho}^{-1} B_{\sigma(t)}{}' M_z(x(t)) \widehat{P}_{\sigma(t)}(t) z(x(t)) \qquad (17)$$

At this point, one should iteratively approach the nearly-optimal solution $P^\circ{}_\sigma$ by updating the parameterization $\widehat{P}_\sigma$ at every time step. The R-PCAO algorithm is a specific way of updating $\widehat{P}_\sigma$ through the use of the building model. The R-PCAO algorithm is schematically represented in Fig. 4, and described through the following steps.

(1) *Calculation of close-to-optimality index*: Select a sampling time dt, and consider the index

$$\epsilon_{\sigma(t)}(x(t), d(t), \widehat{P}_{\sigma(t)}(t)) = \widehat{V}_{\sigma(t)}(t) - \widehat{V}_{\sigma(t-dt)}(t-dt)$$
$$+ \int_{t-dt}^t [H_{\sigma(t)}(t) + \bar{\rho} u'_{\sigma(t)}(t) u_{\sigma(t)}(t)] dt \qquad (18)$$
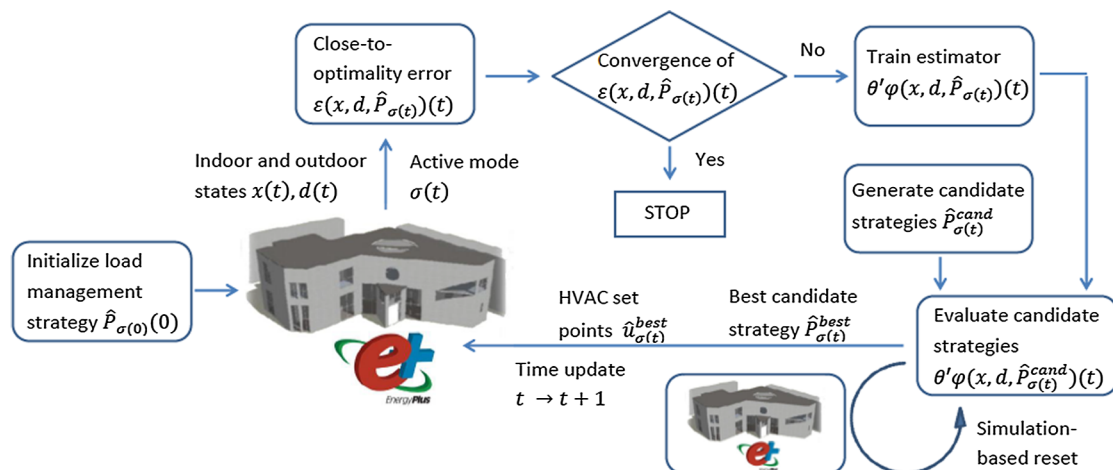

Fig. 4. R-PCAO load management and zoning strategy.

which can be referred to as the close-to-optimality index. In fact, the smaller the $E(\widehat{P}_\sigma) = \epsilon^2$, the closer $\widehat{P}_\sigma$ is to P°_σ (albeit the approximation error $O(1/L)$). Therefore, one can think about using a gradient-like descent for updating $\widehat{P}_\sigma$,

$$\widehat{P}_{\sigma(t+dt)}(t+dt) = \widehat{P}_{\sigma(t)}(t) - a(t)\nabla_{\widehat{P}}\ \epsilon^2(\widehat{P}_{\sigma(t)}(t)) \tag{19}$$

where $a(t) > 0$ is an update step, so as to minimize ϵ^2. In fact, when $\widehat{P}_\sigma$ converges to the optimal P°_σ, then ϵ^2 would be of the order of the approximation error for all time steps. However, (19) cannot be directly used because:

a. An analytic expression of $\nabla_{\widehat{P}} E$ in (19) is not available. In fact, $\nabla_{\widehat{P}} E$ depends on the simplified dynamics (11), which are an approximation of the actual building dynamics;

b. It is well know that, in practice, the approximation error term $O(1/L)$ can make the convergence properties of the standard gradient descent algorithm invalid [41].

To overcome these technical difficulties we construct an alternative descent method according to the following steps.

(2) *Update linear-in-the-parameters estimator*: Consider the linear-in-the-parameters estimator

$$\widehat{\epsilon}_{\sigma(t)}(x(t), d(t), \widehat{P}_{\sigma(t)}(t)) = \theta'\phi(\widehat{P}_{\sigma(t)}(t)) \tag{20}$$

$$\theta = \arg\min_{\vartheta} \sum_{l=t-T}^{k} (\widehat{\epsilon}_{\sigma(t)}(\widehat{P}_{\sigma(l)}(l)) - \vartheta'\phi(\widehat{P}_{\sigma(l)}(l)))^2$$

to approximate $\nabla_{\widehat{P}}\ \epsilon^2$, i.e. to approximate the gradient of the objective function with respect to $\widehat{P}_{\sigma(t)}$. Using stochastic approximation techniques, it has been shown in [41] that such an approximation will iteratively converge close to the actual gradient as more data are collected.

(3) *Generate candidate strategies*: Because $\widehat{P}_\sigma$ is the parameterization of a Lyapunov function, only positive definite matrices should be considered for every update. This can be easily achieved by generating the appropriate candidate perturbations $\Delta\widehat{P}_{\sigma(t)}^{(i)}$, $i = 1, 2, ..., N$

$$\widehat{P}_{\sigma(t)}^{(i)\ (cand)} = (1-a(t))\widehat{P}_{\sigma(t)}^{best}(t) + a(t)\Delta\widehat{P}_{\sigma(t)}^{(i)} \tag{21}$$

where $a(t)$ is a positive update step and $\widehat{P}_\sigma^{best}$ will be defined later. It is immediate to see that, if $\epsilon_1 I \leqslant \widehat{P}_{\sigma(t)}^{best}(t) \leqslant \epsilon_2 I$ and $\epsilon_1 I \leqslant \Delta\widehat{P}_{\sigma(t)}^{(i)} \leqslant \epsilon_2 I$, then $\epsilon_1 I \leqslant \widehat{P}_{\sigma(t)}^{(i)}(t) \leqslant \epsilon_2 I$ at every time t.

(4) *Evaluate candidate strategies* and *Select best strategy*: Every $\widehat{P}_{\sigma(t)}^{(i)\ (cand)}$ represent a possible load management strategy to the tested. Clearly, it is computationally too demanding to choose the brute force approach of testing each $\widehat{P}_{\sigma(t)}^{(i)\ (cand)}$ on the building model. On the other hand, one can evaluate each candidate perturbations $\widehat{P}_{\sigma(t)}^{(i)\ (cand)}$ via the estimator (20), and only the best one (according to the estimator) can be selected for actual test in the building model

$$\widehat{P}_{\sigma(t)}\left(t+dt\right) = \arg\min_{i=1,...,N}\ \widehat{\epsilon}\left(x(t), d(t), \widehat{P}_{\sigma(t)}^{(i)\ (cand)}\right) \tag{22}$$

Therefore, the use of the estimator (20) has the clear advantage that only one evaluation using the building model is performed for each time step. This is important because for each evaluation of a strategy using a building model (simulation-based evaluation), the computational cost is in general proportional to the simulation horizon that one wants to test.

(5) *Simulation-based reset*: To reduce the degrading effect of the approximation error term $O(1/L)$, let us consider a reset based on the simulation-based evaluation. In other words, one should memorize $\widehat{P}_{\sigma(t)}^{best}(t)$ as the strategy with the best simulation-based performance. At this point, as highlighted in (21), the approximation of the value function $\widehat{P}_{\sigma(t)}(t+dt)$ at the next time step is always calculated starting from a perturbation of $\widehat{P}_{\sigma(t)}^{best}(t)$.

4.3. Switched-based approximation

Any building management system exploits a certain set of information to operate its management actions. Such information can be typically categorized as internal feedback factors (internal temperatures, generally denoted with x) and external measurable feedback factors (external temperature and solar radiation, generally denoted with d). The same holds for R-PCAO, where the information of x and d should be used to operate an optimal management: in particular, in R-PCAO, the measurements are used to approximate the value function and the control law. In many applications, a quadratic approximation of the value function, e.g. $x'\widehat{P}x$ and a linear approximation of the control law, e.g. $\widehat{u} = -\bar{\rho}^{-1}B'\widehat{P}x$ can provide acceptable performance. However, because we have seen in (3)–(5) that switching modes will occur, the dynamics corresponding to each $\sigma(t)$ which be too different to be handled by a single (linear) controller. This implies that the quadratic/linear approximations must be overcome. Therefore, we utilize different controllers depending on the active mode $\sigma(t)$. In this case, $\widehat{P}_{\sigma(t)}$ and $z(x)$ are

$$\widehat{P}_{\sigma(t)}(k) = \begin{bmatrix} P_{\sigma^{(1)}(t)} & 0 & \cdots & 0 \\ 0 & P_{\sigma^{(2)}(t)} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & P_{\sigma^{(N_R)}(t)} \end{bmatrix}$$

$$z(t) = \begin{bmatrix} \begin{bmatrix} d(t) \\ Q_{occ}^{(1)}(t) \\ x^{(1)}(t) \\ x^{(1)}(t) \end{bmatrix} \\ \vdots \\ \begin{bmatrix} d(t) \\ Q_{occ}^{(N_R)}(t) \\ x^{(N_R)}(t) \\ x^{(N_R)}(t) \end{bmatrix} \end{bmatrix}$$

where the following things have to be noticed: the feedback vector contains information about outside weather conditions, people occupying a certain room, inside temperature and temperature of the neighboring rooms. This allows us a special block-structure in $\widehat{P}_{\sigma(t)}$ where each submatrix $P_{\sigma^{(i)}(t)}$ is activated depending on the mode of room i.

5. Results

This section focuses on the simulations for the test case of Section 2. The simulations have been conducted with Matlab R2015b, and Energy-Plus 6.1.0, on a PC with 16 GB RAM and Intel 4770 k. The weather data have been taken from the EnergyPlus database, in particular for Athens, year 2011. It is important to underline that the numerical values of the RBLM have been optimized so as to perform good (with average PPD around 9% for a nominal occupancy schedule, and as small energy consumption as possible) over 7 different sets of 7 days during summer (Greek climate, from mid June to early August 2011). Note that the hottest weeks of the season are included in these sets, with some variability: considering the hottest weeks allows us to maximize the need for zoning in the building.

The average occupancy schedule of the building is given in Table 2. Note that sometimes the building is occupied during the whole day and sometimes only during the morning, creating a variability that makes the control design more challenging (because occupant behavior has to be taken into account). The one reported in Table 2 is the average occupancy schedule over the entire building: the actual occupancy schedule for each room consists of some random perturbation of the average schedule, as reported in Fig. 5.

Table 2
Average occupancy schedule over one week (the numbers in the header refer to the hours of a day).

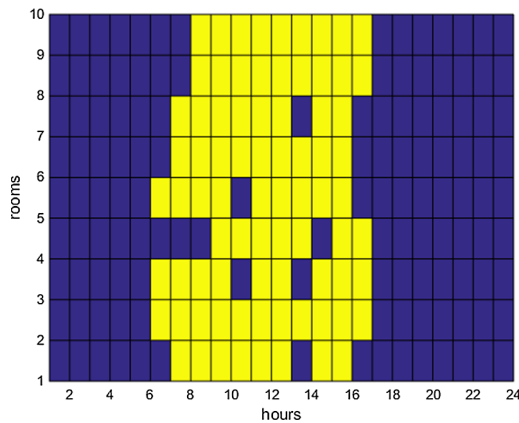|	7–15	8–14 and 17–21	Empty
Monday	✔		
Tuesday		✔	
Wednesday	✔		
Thursday		✔	
Friday		✔	
Saturday	✔		
Sunday			✔

Fig. 5. Example of actual occupancy schedule in the different rooms (dark = unoccupied, light = occupied).

In all subsequent tables we will report the improvement in terms of non-renewable energy consumption (intended as the energy consumed from the grid), and the improvements in terms of PPD. We calculate such costs for the different load management programs, including the proposed one. Furthermore, in order to provide a cumulative (total) improvement, we also calculate

$$\text{Total improvement} = 0.1(\text{Daily energy consumption})[\text{kWh}]$$
$$+ 0.9(\text{Average daily PPD})[\%] \qquad (23)$$

which is a linear combination of the two main performance criteria. The weights 0.1 and 0.9 reflect the same proportion used in the cost matrix H_σ, and thus reflect the actual cost driving the R-PCAO optimization. Different weights would lead to different trade-offs and thus different Pareto-optimal results.

To highlight the benefit of smart zoning, we deploy a 'downgraded' version of R-PCAO where the HVAC set point is the same in all rooms (we call this version R-PCAO no zoning); finally, to prove the necessity of the multi-modal control actions parameterized by the switching signal σ, we also implement a linear version of R-CAO (named R-PCAO linear), which implements zoning, but with a single linear controller (and a single quadratic value function) to handle all possible working modes shown in Fig. 2. The performance of all programs is tested over 7 different sets of 7 days, and a percentage range is given to show the

effect of weather variability on the performance. Table 3 shows the improvement of R-PCAO and RBLM as compared to the two simple FSP strategies. It can be noted that R-PCAO attains total improvements in the range 26–33% as compared to the two FSP strategies. This is a remarkable 39–42% performance improvement over the RBLM strategy (whose total improvement is in the range 15–20% as compared to the two FSP strategies). Such improvements are quite consistent despite different external weather. It is not a surprise that, when no smart zoning action is implemented by R-CAO (R-PCAO no zoning), the performance almost falls back to the performance of RBLM: in fact, RBLM is a strategy optimized via a genetic algorithm that does not implement any smart zoning. Therefore, as compared to optimized strategies with no smart zoning, there is an additional benefit of around 4% reduced energy and 8% improved comfort in implementing a smart zoning program. Finally, the last row of Table 3 reveals that the linear version of R-PCAO (R-CAO linear) attains improvements in the range 21–27% as compared to the two FSP strategies, which is about a 18–19% performance degradation as compared to the proposed (switched) R-PCAO. From here we see that the switching action is fundamental to maximize performance of the load management and zoning program.

For easiness of analysis, let us now focus on one of the 7 sets of Table 3 (in particular on the last set), and let us plot the two bar charts presented in Fig. 6: the first bar chart reports the energy cost (in kWh), while the second bar chart reports the PPD (in %) for all strategies. The important observation is that R-PCAO scores better in both bar charts than all other strategies: for example, as compared to $FSP_{24°C}$, R-PCAO saves more than 150 kWh in 7 days (even with 4% better comfort levels). On the other hand, as compared to $FSP_{25°C}$, R-PCAO has only a slightly better energy cost, but a relevant 6% improvement in thermal comfort. This is particularly remarkable because R-PCAO consumes less energy than a strategy that, by not implementing any pre-cooling action and by raising the set point, is saving a large amount of energy. The reason behind this reduction of energy consumption is that R-PCAO manages to exploit more efficiently the renewable energy from the solar panel. For the same experiment as before, Fig. 7 reports that $FSP_{24°C}$ uses 389 kWh from the solar panel (35% of the total building energy), $FSP_{25°C}$ uses 395 kWh from the solar panel (41% of the total building energy), RBLM uses 475 kWh from the solar panel (44% of the total building energy), and R-PCAO uses a remarkable 590 kWh from the solar panel (51% of the total building energy). The interesting observation is that, if we look at the total energy consumption (non-renewable + renewable) for all strategies, we have: $FSP_{24°C}$ needs 1109 kWh, $FSP_{25°C}$ needs 964 kWh, RBLM needs 1079 kWh, and R-PCAO uses 1157 kWh, which is the largest of all. So, the benefit of R-PCAO is exploiting in a smarter way the renewable energy which is free of charge.

Fig. 8 reports, for both RBLM and R-PCAO, the evolution of the set points for 1-day simulation (the different color variations indicate the different set point levels). The figure shows how R-PCAO is able to give sharper feedback-based adjustments as compared to the RBLM basic working point. In addition, while RBLM does not implement any zoning (all zones are optimized at the same set point, indicated on top of Fig. 8(a)), RPCAO is able to deliver different set points for each room (indicated in Fig. 8(b) with different color degradation). The difference in the set points from room to room can be up to 2 °C. We conclude that

Table 3
R-PCAO Total Cost improvements as compared to $FSP_{24°C}$ and $FSP_{25°C}$ (simulations for 7 different sets of 7 days, summer season).

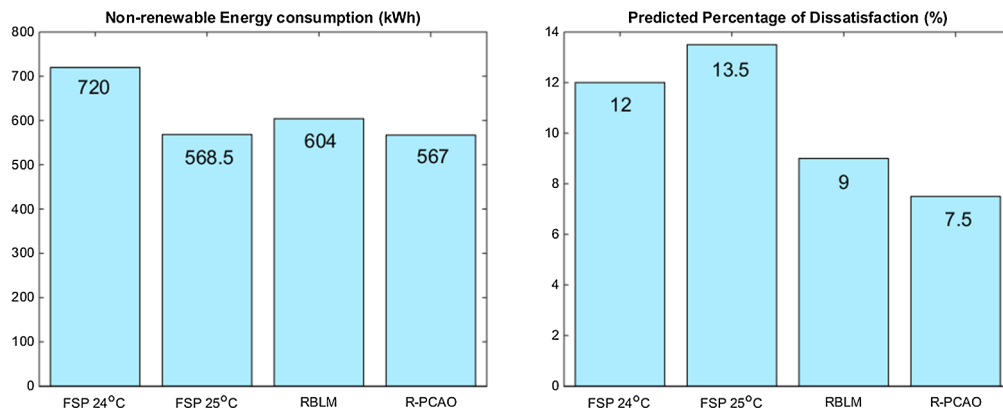|	Energy impr. wrt $FSP_{24°C}$	Comfort impr. wrt $FSP_{24°C}$	Total impr. wrt $FSP_{24°C}$	Energy impr. wrt $FSP_{25°C}$	Comfort impr. wrt $FSP_{25°C}$	Total impr. wrt $FSP_{25°C}$
RBLM	12–14%	24–26%	18–20%	(-8) to (-5)%	32–35%	15–18%
R-PCAO	20–23%	36–39%	30–33%	0–2%	43–47%	26–29%
R-PCAO no zoning	12–14%	25–27%	19–21%	(-7) to (5)%	33–36%	16–19%
R-PCAO linear	16–19%	28–32%	23–27%	(-1) to (1)%	35–38%	21–24%

Fig. 6. Energy and PPD costs for one of the 7 different sets of 7 days.
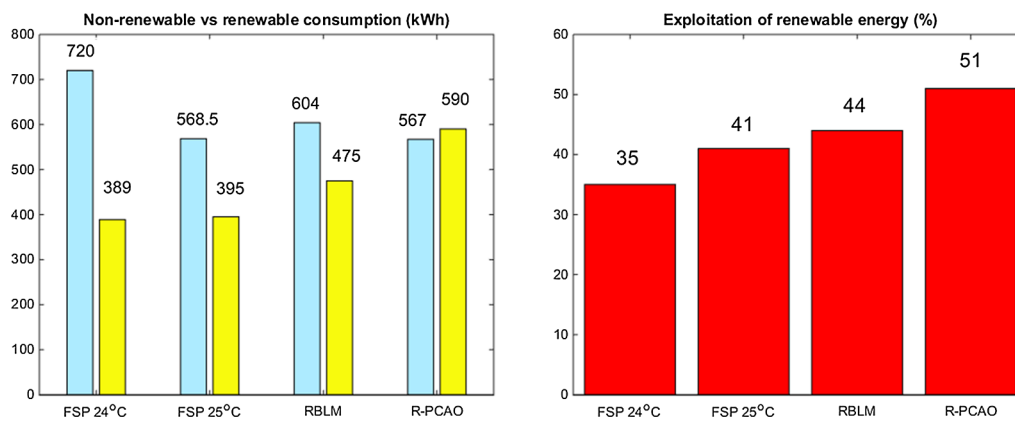
Fig. 7. Exploitation of renewable energy for one of the 7 different sets of 7 days.
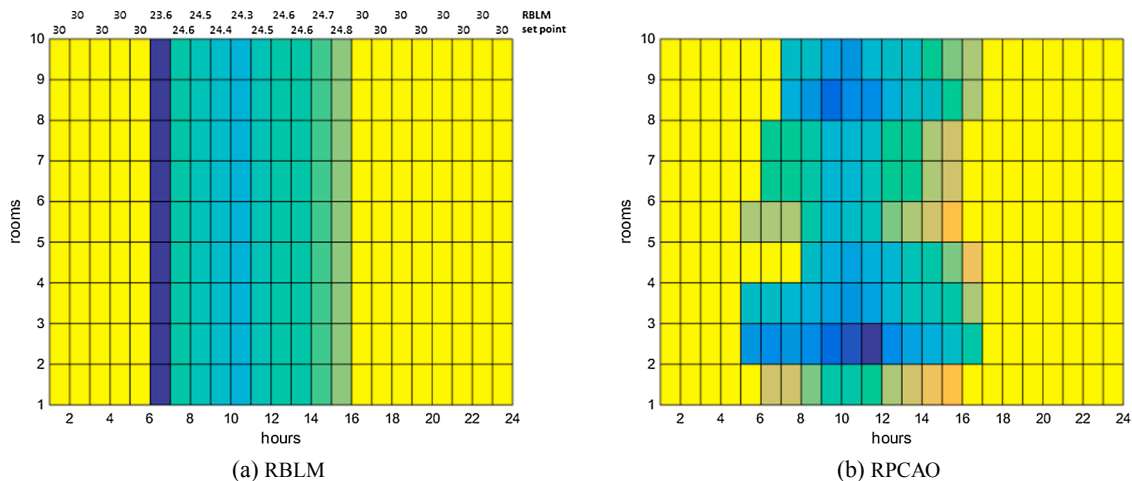

(a) RBLM

(b) RPCAO

Fig. 8. Thermostatic set points during 1-day simulation (the different color variations indicate the different set point levels).

R-PCAO is able to learn the building dynamics from the EnergyPlus data and to use them for feedback-based control.

5.1. Robustness to variable weather

It is clear that rule-based programs can seldom guarantee robustness against weather conditions (unless the rules are weather-dependent, which is not the case for the rules in Section 2.1). In fact, when rule-based programs are optimized over a specific data set, robustness to different data sets may not occur. Additional simulation tests will now demonstrate that the proposed R-PCAO, by providing a feedback-based action, can enhance robustness to different weather conditions. In these

simulations all the strategies are optimized over seven days in early August, and used over 4 different sets of 7 days of early autumn season (mid September-mid October 2011) and over 4 different sets of 7 days of late spring season (mid May-mid June 2011).

The purpose is to test to what extent a load management and zoning program optimized over short data sets can be robust when implemented over longer data sets (this feature is often referred to as *generalization*, i.e. the performance is consistent over data sets different than the data sets used for optimization). Moreover, we perform the optimization of RBLM and R-PCAO during the 4 sets of autumn season and the during the 4 sets of spring season, so as to check the degradation of performance with respect to the best possible performance:

Table 4
Improvements with respect to $FSP_{24°C}$ and $FSP_{25°C}$ (simulations for 4 different sets of 7 days, early autumn season).

	Energy impr. wrt $FSP_{24°C}$	Comfort impr. wrt $FSP_{24°C}$	Total impr. wrt $FSP_{24°C}$	Energy impr. wrt $FSP_{25°C}$	Comfort impr. wrt $FSP_{25°C}$	Total impr. wrt $FSP_{25°C}$
RBLM	9–11%	16–19%	10–12%	(−10) to (−8)%	20–23%	9–10%
R-PCAO	18–21%	32–36%	26–30%	(−1) to (1)%	37–41%	22–25%
RBLM autumn	14–17%	24–27%	17–20%	(−8) to (−6)%	31–34%	14–17%
R-PCAO autumn	19–23%	34–39%	28–33%	0–2%	40–45%	24–28%

Table 5
Improvements with respect to $FSP_{24°C}$ and $FSP_{25°C}$ (simulations for 4 different sets of 7 days, late spring season).

	Energy impr. wrt $FSP_{24°C}$	Comfort impr. wrt $FSP_{24°C}$	Total impr. wrt $FSP_{24°C}$	Energy impr. wrt $FSP_{25°C}$	Comfort impr. wrt $FSP_{25°C}$	Total impr. wrt $FSP_{25°C}$
RBLM	8–10%	15–18%	9–11%	(−11) to (−8)%	18–23%	8–11%
R-PCAO	16–21%	30–34%	24–28%	(−1) to (1)%	39–42%	22–25%
RBLM spring	13–15%	23–27%	16–19%	(−9) to (−5)%	29–34%	13–18%
R-PCAO spring	18–22%	32–36%	27–31%	0–2%	40–44%	23–28%

we refer to such (best possible) strategies as RBLM autumn, R-PCAO autumn, RBLM spring and R-PCAO spring, respectively.

Table 4 presents the results of the early autumn 4-week experiment and Table 5 presents the results of the late spring 4-week experiment. In order to avoid any bias arising from different weather, we compare the performance with respect to the *FSP* strategies. Obviously, the improvements of RBLM and R-PCAO, are worse than the improvements their full data set counterparts RBLM autumn, R-PCAO autumn and RBLM spring, R-PCAO spring, respectively. However, a notable difference arises: while the performance degradation of R-PCAO is acceptable (improvements in the range 21–30% as compared to 24–33% in R-PCAO autumn and 23–31% in R-PCAO spring), the performance degradation for RBLM is higher (improvements in the range 9–12%, as compared to 14–20% in RBLM autumn and 13–19% in RBLM spring. In other words, the average degradation of performance for R-PCAO is below 10%, whereas the average degradation of performance for RBLM is around 40%. Small degradation of performance basically indicates robustness to variability of external condition. Poor robustness of RBLM can be explained by the fact that RBLM does not employ any feedback control: therefore, in order to keep consistent performance under different dynamics and weather conditions, continuous tuning and redesign of RBLM from the user or from the building manager are constantly required. On the other hand, the feedback action embedded in R-PCAO (note that the feedback vector of R-PCAO includes external weather conditions) makes it intrinsically more robust.

5.2. Robustness to variable user behavior (occupancy schedule)

If weather conditions can be imagined to follow some continuous stochastic process, user behavior (in terms of occupancy schedule) will typically follow a discrete stochastic process (e.g. in the form of a Markov chain). Therefore, assessing robustness with respect to variable occupant behavior is even more crucial than with weather conditions. It is clear that the rules in rule-based programs as presented in Section 2.1 are occupancy-dependent. However, because a careful tuning of the set

points in different times of the day and over different rooms is necessary (cf. Fig. 8), robustness is not guaranteed. Therefore, set points optimized over a short data set with a specific occupancy schedule, may not generalize to longer data sets with variable occupancy behavior. Similarly to the previous set of simulations, RBLM and R-PCAO are optimized over seven days in early August, and used over 4 different sets of 7 days. During this 4 sets, we keep the same weather conditions, and we only change the occupancy schedule: the different occupancy schedules are generated as a perturbation of the nominal schedule in Table 2. Moreover, we perform the optimization of RBLM and R-PCAO during the 4 sets of variable occupancy schedules (full data set), so as to check the degradation of performance with respect to best possible performance: we refer to such (best possible) strategies as RBLM schedule and R-PCAO schedule, respectively.

Table 6 presents the results where, in order to avoid any bias arising from different occupancy schedules, we compare the performance with respect to the *FSP* strategies. Because the same weather conditions are used for all weeks, the indicated range refers to the different occupancy schedules (differently from Tables 3–5 where the range was arising because of different weather). By looking at Table 6, a notable difference arises: when optimized over the full data set, the range of improvements of RBLM schedule and R-PCAO schedule is quite narrow (the differences are of the order of 1–2%). On the other hand, optimization over a shorter horizon leads to ranges of the order of 3–4% for R-PCAO and of 6–7% for RBLM, which shows that R-PCAO is more robust to variability in the occupancy schedule. This happens because the feedback action embedded in R-PCAO exploits the relations between the occupancy status and the thermal state of the building, so that such relations can be generalized when different occupancy schedules occur.

We conclude that the proposed R-PCAO program, due to its embedded feedback action (with information of thermal and occupancy conditions), leads to consistent improvements even under changing conditions, covering both variable occupancy patterns and variable weather.

Table 6
Improvements with respect to $FSP_{24°C}$ and $FSP_{25°C}$ (simulations for 4 different sets of 7 days, variable occupancy).

	Energy impr. wrt $FSP_{24°C}$	Comfort impr. wrt $FSP_{24°C}$	Total impr. wrt $FSP_{24°C}$	Energy impr. wrt $FSP_{25°C}$	Comfort impr. wrt $FSP_{25°C}$	Total impr. wrt $FSP_{25°C}$
RBLM	10–14%	17–26%	11–18%	(−10) to (−6)%	23–32%	9–15%
R-PCAO	19–21%	32–36%	27–31%	0–2%	41–43%	24–27%
RBLM schedule	14–15%	25–27%	18–19%	(−7) to (−5)%	32–33%	15–16%
R-PCAO schedule	20–22%	34–36%	29–31%	0–2%	43–44%	26–27%

6. Conclusions and future work

Load management actions in large buildings often neglect the occupant-building dynamic interaction and prevent smart zoning strategies, i.e. setting the temperature in every room at different levels according to geometry, orientation and interaction among rooms caused by variable occupancy patterns. In this work we created a novel self-tuning demand management architecture that adds, on top of a rule-based load management, sharper feedback-based zoning actions: this was achieved by embedding multi-mode (switched) behavior into the approximate solution of a Hamilton-Jacobi-Bellman framework. We demonstrated the proposed load management and zoning program via a test case for intelligent management of heating, ventilating and air conditioning (HVAC) in a building with multiple zones. In particular, to demonstrate the proposed approach, we employ a realistic EnergyPlus model of an actual office building in Crete, Greece. Extensive tests show that the proposed solution is able to learn the building dynamics and to provide different set points in every room in such a way to optimize the whole building performance (exploitation of renewable energy sources with improved thermal comfort). As compared to pre-programmed (non-optimal) strategies, we show that smart zoning makes it is possible to save more than 15% energy consumption, while the thermal comfort results increased more than 25%. As compared to optimized strategies in which smart zoning is not implemented, smart zoning leads to additional 4% reduced energy and 8% improved comfort, demonstrating improved occupant-building interaction. Moreover, the proposed solution is robust to occupant behavior and weather conditions the closed-loop feature of the approach makes it robust to variable weather conditions and occupancy schedules.

Despite its positive performance, this approach is open to future improvement. First, it would be really relevant to include the occupant interaction with the building automation system, for example opening window/doors or overruling the set point: unfortunately, we could not find any 'human model' that could be included in EnergyPlus to simulate such human response: we would like to study this relevant point on future work. Second, load management and zoning could be connected with demand response programs defined by the grid operator via price-based incentives. This scenario could be handled via game-theoretic approaches, to be handled for example in an extended Hamilton-Jacobi-Isaac formulation: studies in this direction can be found in [42]. Finally, some further steps are still necessary to reach a fully-automated program, for example plug-n-play integration of new loads, storage and generation devices. Studies in this direction can be found in [43] with multiple control strategies taking different decisions depending on the available equipment.

Acknowledgment

References

[1] Samad T, Koch E, Stluka P. Automated demand response for smart buildings and microgrids: the state of the practice and research challenges. Proc IEEE 2016;104:726–44.

[2] Abdulaal A, Moghaddass R, Asfour S. Two-stage discrete-continuous multi-objective load optimization: an industrial consumer utility approach to demand response. Appl Energy 2017;206:206–21.

[3] Samad T, Annaswamy AM. Controls for smart grids: architectures and applications. Proc IEEE 2017;105(11):2244–61.

[4] Pérez-Lombard L, Ortiz J, Pout C. A review on buildings energy consumption information. Energy Build 2008;40(3):394–8.

[5] Mohsenian-Rad A-H, Wong VWS, Jatskevich J, Schober R, Leon-Garcia A. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. IEEE Trans Smart Grid 2010;1(3):320–31.

[6] Kie PLT, Theng LB. Intelligent control of heating, ventilating and air conditioning systems. In: Köppen M, Kasabov N, Coghill G, editors, Advances in neuro-information processing, Berlin, Heidelberg; 2009. Berlin Heidelberg: Springer, p. 927–34.

[7] Alcala R, Alcala-Fdez J, Herrera F. A proposal for the genetic lateral tuning of linguistic fuzzy systems and its interaction with rule selection. IEEE Trans Fuzzy Syst 2007;15(4):616–35.

[8] Wang Z, Wang L, Dounis AI, Yang R. Multi-agent control system with information fusion based comfort model for smart buildings. Appl Energy 2012;99:247–54.

[9] Yang L, Yan H, Lam JC. Thermal comfort and building energy consumption implications–a review. Appl Energy 2014;115:164–73.

[10] ASHRAE, ANSI/ASHRAE standard 55-2004: thermal environmental conditions for human occupancy, American Society of Heating, Refrigerating and air-Conditioning Engineers; 2004.

[11] Wu W, Skye HM, Domanski PA. Selecting hvac systems to achieve comfortable and cost-effective residential net-zero energy buildings. Appl Energy 2018;212:577–91.

[12] Mathews EH, Arndt DC, Piani CB, Van Heerden E. Developing cost efficient control strategies to ensure optimal energy use and sufficient indoor comfort. Appl Energy 2000;66(2):135–59.

[13] Alcala R, Casillas J, Cordon O, Gonzalez A, Herrera F. A genetic rule weighting and selection process for fuzzy control of heating, ventilating and air conditioning systems. Eng Appl Artif Intell 2005;18(3):279–96.

[14] Tzivanidis C, Antonopoulos KA, Gioti F. Numerical simulation of cooling energy consumption in connection with thermostat operation mode and comfort requirements for the athens buildings. Appl Energy 2011;88(8):2871–84.

[15] Wang Z, Yang R, Wang L, Green RC, Dounis AI. A fuzzy adaptive comfort temperature model with grey predictor for multi-agent control system of smart building. In: 2011 IEEE Congress on Evolutionary Computation (CEC); 2011. p. 728–35.

[16] Zhou H, Rao M, Chuang KT. Knowledge-based automation for energy conservation and indoor air quality control in hvac processes. Eng Appl Artif Intell 1993;6(2):131–44.

[17] Mossolly M, Ghali K, Ghaddar N. Optimal control strategy for a multi-zone air conditioning system using a genetic algorithm. Energy 2009;34(1):58–66.

[18] Ogunjuyigbe ASO, Ayodele TR, Akinola OA. User satisfaction-induced demand side load management in residential buildings with user budget constraint. Appl Energy 2017;187:352–66.

[19] Michailidis IT, Baldi S, Pichler MF, Kosmatopoulos EB, Santiago JR. Proactive control for solar energy exploitation: a german high-inertia building case study. Appl Energy 2015;155:409–20.

[20] Solano JC, Olivieri L, Caamao-Martin E. Assessing the potential of pv hybrid systems to cover hvac loads in a grid-connected residential building through intelligent control. Appl Energy 2017;206:249–66.

[21] Alibabaei N, Fung AS, Raahemifar K, Moghimi A. Effects of intelligent strategy planning models on residential hvac system energy demand and cost during the heating and cooling seasons. Appl Energy 2017;185:29–43.

[22] Adhikari R, Pipattanasomporn M, Rahman S. An algorithm for optimal management of aggregated hvac power demand using smart thermostats. Appl Energy 2018;217:166–77.

[23] Michailidis IT, Schild T, Sangi R, Michailidis P, Korkas C, Futterer J, Muller D, Kosmatopoulos EB, et al. Energy-efficient hvac management using cooperative, self-trained, control agents: a real-life german building case study. Appl Energy 2018;211:113–25.

[24] Xu X, Culligan PJ, Taylor JE. Energy saving alignment strategy: achieving energy efficiency in urban buildings by matching occupant temperature preferences with a buildings indoor thermal environment. Appl Energy 2014;123:209–19.

[25] Abrol S, Mehmani A, Kerman M, Meinrenken CJ, Culligan PJ. Data-enabled building energy savings (d-e bes). Proc IEEE 2018;106(4):661–79.

[26] Wang W, Chen J, Huang G, Lu Y. Energy efficient hvac control for an ips-enabled large space in commercial buildings through dynamic spatial occupancy distribution. Appl Energy 2017;207:305–23.

[27] Oldewurtel F, Sturzenegger D, Morari M. Importance of occupancy information for building climate control. Appl Energy 2013;101:521–32.

[28] Goyal S, Barooah P, Middelkoop T. Experimental study of occupancy-based control of hvac zones. Appl Energy 2015;140:75–84.

[29] Kim JS, Edgar TF. Optimal scheduling of combined heat and power plants using mixed-integer nonlinear programming. Energy 2014;77:675–90.

[30] Verhelst C, Logist F, Van Impe J, Helsen L. Study of the optimal control problem formulation for modulating air-to-water heat pumps connected to a residential floor heating system. Energy Build 2012;45:43–53.

[31] Korkas CD, Baldi S, Michailidis I, Kosmatopoulos EB. Intelligent energy and thermal comfort management in grid-connected microgrids with heterogeneous occupancy schedule. Appl Energy 2015;149:194–203.

[32] Korkas CD, Baldi S, Michailidis I, Kosmatopoulos EB. Occupancy-based demand response and thermal comfort optimization in microgrids with renewable energy sources and energy storage. Appl Energy 2016;163:93–104.

[33] Ozturk Y, Senthilkumar D, l Kumar S, Lee G. An intelligent home energy management system to improve demand response. IEEE Trans Smart Grid 2013;4(2):694–701.

[34] Roy S, Mishra K, Maulik U, Basu S. A distributed multilabel classification approach towards mining appliance usage in smart homes. In: 2017 IEEE Calcutta Conference (CALCON); 2017. p. 474–8.

[35] De Silva LC, Morikawa C, Petra IM. State of the art of smart homes. Eng Appl Artif Intell 2012;25(7):1313–21.

[36] Nguyen TA, Aiello M. Energy intelligent buildings based on user activity: a survey. Energy Build 2013;56:244–57.

[37] Ahmadi-Karvigh S, Ghahramani A, Becerik-Gerber B, Soibelman L. Real-time activity recognition for energy efficiency in buildings. Appl Energy 2018;211:146–60.

[38] Crawley DB, Hand JW, Kummert M, Griffith BT. Contrasting the capabilities of building energy performance simulation programs. Build Environ 2008;43(4):661–73.

[39] Crawley DB, Lawrie LK, Winkelmann FC, Buhl WF, Huang YJ, Pedersen CO, Strand RK, Liesen RJ, Fisher DE, Witte MJ, et al. Energyplus: creating a new-generation building energy simulation program. Energy Build 2001;33(4):319–31.

[40] Kosmatopoulos EB, AGILE consortium. AGILE project, supported by the european commission fp7-ict-5-3.5, under the contract #257806; 2018. <https://cordis.europa.eu/project/rcn/95306_en.html>.

[41] Baldi S, Michailidis I, Kosmatopoulos EB, Ioannou PA. A plug and play computationally efficient approach for control design of large-scale nonlinear systems using cosimulation: a combination of two ingredients. IEEE Control Syst 2014;34(5):56–71.

[42] Yu M, Hong SH. A real-time demand-response algorithm for smart grids: a stackelberg game approach. IEEE Trans Smart Grid 2016;7(2):879–88.

[43] Baldi S, Michailidis I, Ravanis C, Kosmatopoulos EB. Model-based and model-free plug-and-play building energy efficient control. Appl Energy 2015;154:829–41.