Efficient HVAC controls: A symbolic approach

Ondrej Holub, Majid Zamani, and Alessandro Abate

Abstract— **Heating, ventilation, and air conditioning (HVAC) systems contribute significantly to energy consumption in buildings, and are the target of an ongoing push towards energy efficiency, which reflects on new requirements for the configuration and operation of HVAC units. On-off actuators are being replaced by staged or modulated ones, and additional components are being added to the HVAC units, which have to be properly managed by new, advanced control architectures. As such, it is increasingly more difficult for traditional control approaches to cope with new devices and meet new objectives. This leads to an opportunity for alternative, novel approaches to be adopted by the industry: promising are recently developed formal methods for controller synthesis. In this article, we employ formal, symbolic techniques to synthesize controllers for a model of a roof-top unit conditioning a low-rise commercial building. The synthesized controller is validated on a simulation with practically relevant operating conditions.**

I. Introduction

Reduced greenhouse gas emissions and decreased dependence on non-renewable energy resources are commonly recognised as key societal challenges. Authorities in many countries have been introducing regulations to stimulate and enforce measures leading to reduced energy consumption [11]. Buildings consume a significant amount of energy: it is estimated they contribute to 20-40% of total consumption in developed countries [16]. As such, they have been targeted by the regulations too.

Heating, ventilation, and air conditioning (HVAC) systems are the single largest contributor to energy consumption in buildings. While building energy efficiency regulations differ over individual countries, their impact on HVAC systems is in general two-fold: they impose changes either in the configuration or in the operation of building automation systems. Typical examples are on-off actuators, which are being replaced by staged or modulated ones to improve partial-load efficiency of the HVAC system (efficiency at a fraction of the rated ventilation and heating/cooling loads). Such retrofits include multi-speed fans [9] for cost-sensitive applications or variable-speed compressors [10] for high-performance systems. Additional components such as economising and heat recovery units are being added to HVAC systems to optimise energy flows therein [20]. As a result, new interactions between extended components and enhanced subsystems need to be taken into account by HVAC control systems, in order to achieve the targeted efficiency improvements. Traditional control approaches, which combine simple PI loops via rule-based reset logic, are becoming too complex to commission or not suitable at all due to their lack of flexibility. New controller synthesis approaches are thus needed, which can facilitate the transition from rule-based, on-off controls to energy-efficient complex strategies.

Opportunities for new control strategies

In the context of modern energy-efficient HVAC retrofits, the following opportunities can be defined for the synthesis of new control strategies.

Let us first discuss the controller design for existing HVAC systems. The replacement of an existing control strategy is usually significantly cheaper than the retrofitting of HVAC equipment. Further, many HVAC systems are operated inefficiently: examples of inefficiencies include oscillating loops leading to an increase in equipment wear, or wrong logical control rules for economising components such as outside air dampers. As such, there is an opportunity for better controller synthesis schemes to serve as a *low-cost energy efficiency measure*.

Let us now elaborate on the controller design for retrofitted HVAC systems. Whenever a decision is made to invest in new HVAC components, the control strategy needs an update, since retrofitting often renders the original control strategy outdated. This is for instance the case with the replacement of on-off components with staged or modulating ones. Notice that return of investment estimates do not provide much insight over the implemented control structures, because such calculations are based on the assumption of perfect controllers: new control strategies are instead expected to *deliver a performance improvement*.

Finally, with the goal of *evaluating the savings potential*, the consumption of existing HVAC systems can be compared with simulated performance of various retrofit options. Calculations involving realistic performance of the closed-loop systems based on a unified methodology for controller synthesis would improve the reliability of such assessment of retrofit options.

Key aspects for quality of controllers

We identify the following aspects to be critical for the deployment of the control strategies in the applications discussed above:

- HVAC units operate in variable conditions (such as ambient temperature and humidity, and temperature, humidity and quality of conditioned air), and as such their controllers need to satisfy performance specifications under all the situations.
- HVAC systems consist of many interacting subsystems: different combinations of their components are activated depending on the operating conditions. Adequate approaches to controller synthesis are needed to handle such interactions.

This work has been supported by the EC IAPP project AMBI 324432.

O. Holub is with ACS Global Labs, Honeywell, Prague, Czech Republic ondrej.holub@honeywell.com

M. Zamani is with the Department of Electrical and Computer Engineering, Technische Universität München, D-80290 Munich, Germany. zamani@tum.de

A. Abate is with the Department of Computer Science, University of Oxford, OX1 3QD, Oxford, United Kingdom. alessandro.abate@cs.ox.ac.uk

- The behaviour of HVAC units depends on a large number of parameters, and the information available for controller synthesis and for operation is limited and uncertain.
- The health of the HVAC system is key to its energy consumption: any stress conditions, such as cycling or overloading, needs to be avoided to minimise wear and tear of the system.
- HVAC controls are specified, implemented, and installed as part of larger design or retrofit projects, which are subject to stringent constraints. Commissioning of the employed controllers should require only minimum efforts and skills.

This work puts forward the use of new control architectures that have the potential to target the aspects above.

Related work

Existing work targets aspects of the discussed high-level requirements. HVAC systems have been a traditional testbed for contemporary popular controller synthesis methods, based on neural networks [7], LQG control [25], $\mathcal{H}_\infty$ methods [1], [23], and model predictive control (MPC) [6].

Despite all these efforts, there is still a gap between academic research and industrial practice. The work in [14] has discussed practical limits of the existing architectures, all of which share insufficient analysis of the cost-to-benefit ratio: whilst most of the published methods focus on improved performance, in practice it appears that the improvement is not relevant enough, the sustainability of the performance in real operating conditions is uncertain, and the commissioning of the existing strategies is too costly or unfeasible.

An exception is MPC applied to real-time optimisation of system set-points [4], [21], [22]; however, these contributions exclusively focus on set-point regulation and assume that the plant controllers are able to follow them. General application of MPC to the synthesis of low-level controllers [6] is in practice questionable due to the commissioning complexity, the small inertia of most HVAC units, and the largely unpredictable operating conditions.

Contributions of this work

To the best of the authors' knowledge, this paper is the first to apply symbolic formal methods to industrially relevant controller synthesis, in particular for a model of HVAC roof-top units. These methods can be used in situations encountered by energy service companies in the iterative process of bidding and of implementation of HVAC retrofits [12]. The key advantages of symbolic formal methods are the *flexibility* to deal with various types of controlled systems (legacy as well as modern HVAC units) and the possibility to provide *formal guarantees*. Such a combination is fairly unique and as such offers differentiating features, as discussed in more detail in Section III. Towards other goals, symbolic techniques based on formal abstractions have been deployed on models for smart buildings in [19], [18].

The mentioned requirements on the controller architectures are addressed by the synthesis problem studied in this article as follows. The limited information about the controlled system is reflected by using a simple and conservative system description for controller synthesis. The variability of conditions are addressed by synthesising dedicated strategies for each major type of operating conditions. The interactions between subsystems are handled in an optimal way thanks to the multivariable nature of the synthesis method. Any stress is removed from the system operation by designing a controller whose outputs are modified only if the previously used ones cannot be employed again. Finally, the synthesis aims at a control strategy which would be pre-engineered as a part of the controller's firmware. Hence the only configuration to be done during deployment and commissioning of the controller consists in choosing the right strategy from a library based on the controlled and manipulated variables and their properties.

II. Description of the Case Study

Setup and control objectives

We have selected a simple HVAC system for demonstration, comprising a packaged direct expansion rooftop unit (RTU) that conditions one zone in a single storey building. As shown in Figure 1, the building is split in six zones, each conditioned by a dedicated RTU. Even though these zones interact with each other, the impact is marginal during normal operation[1] and as such is not considered when designing the controller for a single zone.

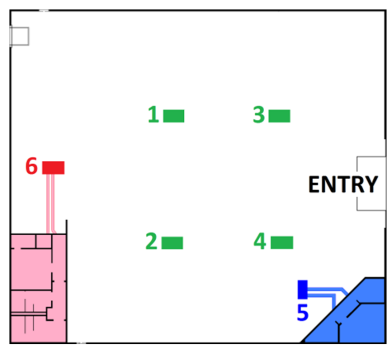

Fig. 1. Floor plan of the considered commercial building: RTUs (points 1-4) condition open-space zones in the sales area, whereas smaller RTUs are dedicated to manager's office (point 5) and facilities (point 6).

Direct-expansion RTUs (Figure 2) use a refrigerant vapour expansion-compression cycle to directly cool the supplied air [8]. Their cooling power is controlled by the number of compressors used in the cycle: two- to four-compressor units are the most common. A supply fan blows the air across the evaporator, which serves as a cooling coil. In the simplest setting, the supply air is directly transported to the conditioned space. Economiser dampers can be used to mix fresh outside air with air returned from the zone (Figure 2) to alter the air properties at the intake of the cooling/heating coils. The considered RTU is equipped with a two-stage compressor, a multi-speed supply fan and a modulating economiser. For the sake of simplicity, the economiser is considered fixed (which means that all the three dampers shown in Figure 2 remain in constant positions).

Each RTU is wired to a plant controller. The objective of the RTU controller is to regulate heat and moisture exchange to the space in order to meet a temperature set point and to keep the space humidity within an acceptable range. For simplicity, in this work only a cooling RTU is considered (that is, heating is not available within the unit). Typical

[1] Interaction effects can be added to the internal gains in (II.1) and (II.2).

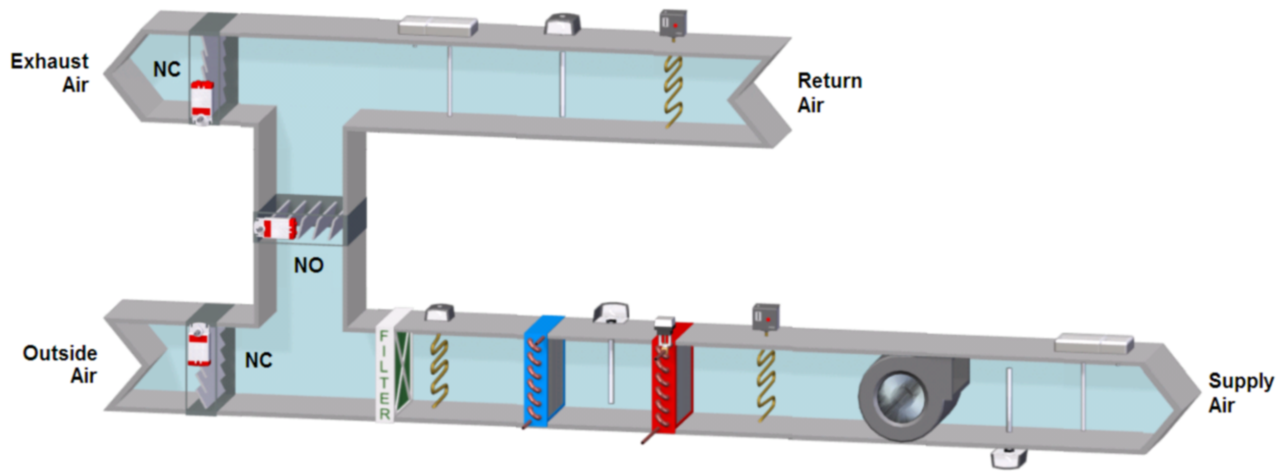

Fig. 2. Roof-top unit. Direct-expansion cooling (blue, left), heating coil (red, right), supply fan and economiser dampers (NO and NC refer to the default damper positions: open and close, respectively). Further in the picture are indicated the locations of temperature and humidity sensors of return, mixed and supply air – these sensors are not employed in this work.

TABLE I
SPECIFICATIONS OF RTU AND PROPERTIES OF CONTROL SIGNALS

Output variables (measured)

Name	Target	Range	Resolution
Zone temp.	$21 \pm 1°\mathrm{C}$	$11 - 31°\mathrm{C}$	$0.1°\mathrm{C}$
Rel. humid.	$30 - 70\%$	$10 - 90\%$	1%

Input variables (actuated)

Name	Stages	Range	Values
Supply fan	4	$0 - 100\%$	$0, 25, \dots, 100$
Compressor	2	$0 - 100\%$	$0, 50, 100$

specifications are shown in Table I: $\pm 1°\mathrm{C}$ band around $21°\mathrm{C}$ temperature set point, and 30% to 70% range for the humidity, which is considered secondary objective (i.e. it should not let the zone temperature go out of range). The only sensors available are positioned in the zone and accessed by the thermostat: temperature and relative humidity of the conditioned zone air are measured with $0.1°\mathrm{C}$ and 1% resolution, respectively.

The dynamics of the control model are described next.

Model of the controlled RTU system

As will be discussed below, an approximate linear dynamic model will be used for controller synthesis. Before that, the RTU system is introduced by qualitative description based on first principles [5], [24].

The heat and moisture gains in the zone are compensated by the supplied air as follows:

$$\begin{aligned} C_{\mathrm{ZA}}\dot{T}_{\mathrm{ZA}} &= P_{\mathrm{ZA}} - P_{\mathrm{RTU}} && \text{(II.1)} \\ &= P_{\mathrm{ZA}} - C_{\mathrm{AIR}}\dot{m}_{\mathrm{SA}}(T_{\mathrm{ZA}} - T_{\mathrm{SA}}) \\ \dot{w}_{\mathrm{ZA}} &= h_{\mathrm{ZA}} - h_{\mathrm{RTU}} && \text{(II.2)} \\ &= h_{\mathrm{ZA}} - \dot{m}_{\mathrm{SA}}(w_{\mathrm{ZA}} - w_{\mathrm{SA}}), \end{aligned}$$

where P, h, $\dot{m}$, T, and w represent the heat gains, moisture gains, mass flow, temperature, and absolute humidity, respectively; and the subscripts $(.)_{\mathrm{ZA}}$ and $(.)_{\mathrm{SA}}$ refer to zone and supply air, respectively.

The amount of heat P_{RTU} extracted from the supply air by the cooling coil depends in general on the compressor power, mixed air temperature and humidity (T_{MA} and w_{MA}, resp.), outside air temperature T_{OA}, and supply air flow $\dot{m}_{\mathrm{SA}}$. The properties of the mixed air at the intake of the cooling coil depend on the outside air and return air properties, and the mixing ratio for the two air streams. The mixing ratio is determined by the economiser's position. The airflow is a function of the fan speed and system resistance, dominated by the ductwork and the air filter (Figure 2). The resistance to airflow is further altered by the economiser positions and by the pressure changes in the conditioned space (this can be caused for example by opening the windows).

The RTU performance is usually presented via a datasheet made from data that is collected in a laboratory setup for specific operating conditions only. Performance under real operating conditions is subject to large uncertainty introduced by various sources, including: properties of the building as well as the supply/return duct, location and neighbourhood of the building (weather and shading), location and quality of the sensors, occupancy and usage of the building, and interaction between conditioned zones. RTU controllers are traditionally preconfigured and have only a limited amount of adjustability, which allows contractors to change preconfigured settings in a number of steps.

While of general interest the first-principle equations (II.1)-(II.2) cannot be directly used for controller synthesis, and are thus replaced by a simpler approximate model. Specifically, a linear approximation is used to describe the local behaviour of the plant, around operating points described by the zone set points, fixed heating and moisture loads and the calculated RTU actions. The following operating conditions are important for the RTU-based system above: a *nominal case* that represents a typical cooling load for which the RTU plant was specified; a *low cooling demand* common in the low season, during which it is desired to keep the switching between stages at reasonably low frequency; a *low recirculation* of air (economiser widely open), during which the impact of the compressor and of the fan heavily depend on outside air properties; and *high humidity* conditions which lower the effect of the compressor on the temperature.

A state-space representation of the linearised plant model Σ can be constructed as follows:

$$\Sigma \left\{ \begin{aligned} \dot{\xi} &= A\xi + B(\upsilon + \omega), \\ \zeta &= C\xi, \end{aligned} \right. \qquad \text{(II.3)}$$

where

$$A = \begin{bmatrix} A_1 & 0_{2\times 2} \\ 0_{2\times 2} & A_2 \end{bmatrix}, B = \begin{bmatrix} B_1 \\ B_2 \end{bmatrix}, C = \begin{bmatrix} C_1 \\ C_2 \end{bmatrix}. \tag{II.4}$$

The input and output variables follow Table I: υ_1 is supply fan speed, υ_2 represents the compressor stages used, ζ_1 and ζ_2 are zone temperature and relative humidity, respectively. Each input-output pair is represented by a first-order low-pass filter. The first and third elements of the state vector ξ correspond to the temperature and humidity, respectively: hence $C_1 = [1\ 0\ 0\ 0]$, $C_2 = [0\ 0\ 1\ 0]$.

The input and state matrices, on the other hand, depend on static gains and time constants of the individual input-output channels and vary with operating conditions. In the nominal operating condition, the roof-top unit selected for validation has the following state-space description[2]:

$$A_1 = 10^{-4}\begin{bmatrix} -28 & -5.6 \\ 0 & -8.3 \end{bmatrix}, A_2 = 10^{-4}\begin{bmatrix} -17 & 1.0 \\ 0 & -2.8 \end{bmatrix}$$
$$B_1 = 10^{-4}\begin{bmatrix} -0.8 & -1.7 \\ 0 & 5.8 \end{bmatrix}, B_2 = 10^{-4}\begin{bmatrix} -1.7 & 0.08 \\ 0 & 2.3 \end{bmatrix}. \tag{II.5}$$

Additionally, as anticipated above the input ω is introduced as an additive disturbance and is used to capture uncertainty of the model (II.3) around (II.5) and short-term fluctuations in the operating conditions (primarily the loads P_{ZA} and h_{ZA}). The disturbance ω is modelled as a low-pass filtered zero-mean Gaussian process e, satisfying the following differential equation:

$$22 \times 10^6 \ddot{\omega} + 36 \times 10^3 \dot{\omega} + \omega = 1.3 \times 10^3 e. \tag{II.6}$$

Gain and time constants of the filter are tuned to mimic amplitude and rate of change of the heat gains observed in real-world applications. While (II.6) is employed for validation, the controller synthesis procedure uses a simpler and conservative representation of disturbance, considered to be constant (between sample times) with values in a range from -20% to $+20\%$ for the first input entry and from -40% to $+40\%$ for the second input entry.

The goal for the controller synthesis is to attain the specification expressed in Table I. This is achieved by identifying the general operating conditions (primarily dependent on current weather and time of the day) and applying a control law synthesised for the current conditions. Switching between control laws is justified by the slow speed with which the conditions change from one type to another.

The next section describes controller synthesis for one type of the conditions, specifically for the nominal case which corresponds to the use of one compressor ($u_2 = 50\%$) and supply fan running at mid-speed ($u_1 = 50\%$). The dynamic model assumed for controller synthesis is (II.5) and the disturbance is assumed to have the succinct dynamical representation, as explained after equation (II.6). Due to the local linearisation mentioned above, the targets for the controlled variables now become $y_1 = 0°\text{C}$ and $y_2 = 0\%$, with some allowed tolerance. The static parameters of the linearised system used as inputs to the controller synthesis are provided for clarity in Table II.

TABLE II

PROPERTIES OF LINEARIZED CONTROLLER SIGNALS

Output variables (measured)

Name	Target	Range	Resolution
Zone temp.	$\pm 1°$C	$\pm 10°$C	$0.1°$C
Rel. humid.	$\pm 5\%$	$\pm 30\%$	1%

Input variables (actuated)

Name	Stages	Range	Values
Supply fan	4	$\pm 25\%$	$-25, 0, 25, 50$
Compressor	2	$\pm 50\%$	$-50, 0, 50$

III. Controller Synthesis via Formal Methods

Controller synthesis framework

We need to cope formally with quantised time, output measurements, and inputs (as expressed in Table II), as common in many industrial control problems. Further, we are interested to provide a *lazy controller*, i.e., a controller where the input is exclusively modified in time when the previously-used input cannot be employed again, which is highly desirable specially in the HVAC control setup. We leverage *symbolic* techniques to provide a provably-correct controller taking into account the aforementioned conditions. In order to do so, first we need to construct a *symbolic model* of the original system Σ in (II.3). Symbolic models are discrete and finite approximations of the concrete continuous dynamics, and are constructed in such a way that controllers designed for the approximations can be refined into controllers for the concrete, original dynamical models. In this work we leverage the techniques in [29] to construct a symbolic abstraction of the concrete system Σ in (II.3).

Consider a set of continuous states $\mathsf{X} \subseteq \mathbb{R}^n$, a set of control inputs $\mathsf{U} \subseteq \mathbb{R}^m$, and a set of disturbance inputs $\mathsf{D} \subseteq \mathbb{R}^m$: all of them are assumed to be finite unions of boxes. Consider a triple $\mathsf{q} = (\tau, \eta, \mu)$ of quantisation parameters, where τ is the sampling time, $\eta \leq \mathit{span}(\mathsf{X})$ is the state-space quantisation, and $\mu \leq \mathit{span}(\mathsf{U})$ is the input set quantisation. A symbolic model $S_{\mathsf{q}}(\Sigma)$ of Σ is a tuple

$$S_{\mathsf{q}}(\Sigma) = \left(X_{\mathsf{q}}, X_{\mathsf{q}0}, U_{\mathsf{q}}, \xrightarrow[\mathsf{q}]{}, Y_{\mathsf{q}}, H_{\mathsf{q}}\right)$$

consisting of:

- $X_{\mathsf{q}} = [\mathsf{X}]_\eta$;
- $X_{\mathsf{q}0} = [\mathsf{X}]_\eta$;
- $U_{\mathsf{q}} = [\mathsf{U}]_\mu$;
- $x_{\mathsf{q}} \xrightarrow[\mathsf{q}]{u_{\mathsf{q}}} x'_{\mathsf{q}}$ if $\mathcal{B}_\eta(x'_{\mathsf{q}}) \cap \mathbf{Post}_{u_{\mathsf{q}},\mathsf{D}}\left(\mathcal{B}_{\eta+\epsilon}(x_{\mathsf{q}})\right) \neq \varnothing$, where $\mathbf{Post}_{u_{\mathsf{q}},\mathsf{D}}\left(\mathcal{B}_{\eta+\epsilon}(x_{\mathsf{q}})\right)$ denotes an over-approximation of the set $\{\xi_{x\upsilon\omega}(\tau) \mid \forall x \in \mathcal{B}_{\eta+\epsilon}(x_{\mathsf{q}}),\ \upsilon(t) = u_{\mathsf{q}},\ \forall t \in [0,\ \tau[,\ \forall \omega \in \mathcal{D}^3\}$ and $\xi_{x\upsilon\omega}(\tau)$ denotes the value of the state trajectory of the original system Σ at time τ under the control input υ and the disturbance ω from initial condition $\xi_{x\upsilon\omega}(0) = x$;
- $Y_{\mathsf{q}} = \mathsf{X}$;

[2]Model parameters were identified from data collected during system operation.

[3]The set $\mathcal{D}$ denotes the set of all piecewise constant functions of duration τ, and taking values in D.

- $H_{\mathsf{q}} = \imath : X_{\mathsf{q}} \hookrightarrow Y_{\mathsf{q}}$.

The constant $\epsilon \in \mathbb{R}_0^+$ in the definition of $S_{\mathsf{q}}(\Sigma)$ denotes the maximum error between the actual state of the system Σ and the state of the observer, called $\hat{\Sigma}$, at time τ, i.e., $\epsilon = \max_{x,\hat{x}\in\mathsf{X}} \|\xi(\tau) - \hat{\xi}(\tau)\|$, where ξ and $\hat{\xi}$ denote the state trajectory of the system and the observer, respectively, and $x = \xi(0)$ and $\hat{x} = \hat{\xi}(0)$. Note that if sets X and U are bounded, which is always the case in practice, then $S_{\mathsf{q}}(\Sigma)$ is called finite or symbolic. When finite symbolic models exist and can be constructed, we can leverage the apparatus of finite-state reactive synthesis [13] towards the problem of designing hybrid controllers enforcing complex logic specifications on the concrete models. Note that the results in [29] require some completeness assumption on the original concrete systems in order to construct their symbolic models. Linear control systems, e.g. Σ described in (II.3), always satisfy the required completeness assumption.

Specifications

One of the great advantages of using symbolic techniques is the ability of enforcing logical specifications on the original concrete systems which are hard (or even impossible) to enforce with classical control techniques. Examples of those logical specifications include the ones expressed via linear temporal logic (LTL) or as automata on infinite strings [3]. As an example, in the HVAC control problem here (c.f. Table II), we are interested to synthesize a controller enforcing outputs of the system enter the set "Target" in Table II, i.e. $W := [-1\ 1] \times [-5\ 5]$, in finite time and stay there forever: the LTL formula[4] encoding this goal is $\Diamond\Box W$.

Implementation and simulations

We now experimentally demonstrate the effectiveness of the symbolic techniques on the HVAC control problem. In this example, the computation of the abstract system $S_{\mathsf{q}}(\Sigma)$ has been implemented by the software tool Pessoa [15] on an iMac with CPU 3.5GHz Intel Core i7. We have assumed that the control inputs are piecewise constant of duration τ and that they take values in U_{q}. The controller enforcing the specifications of interest has been found by standard algorithms from game theory [13], as implemented in Pessoa. For this system and as mentioned in Table II, we can assume $\mathsf{X} = [-10\ 10] \times [-10\ 10] \times [-30\ 30] \times [-30\ 30]$, $\mathsf{U} = [-25\ 50] \times [-50\ 50]$, $\mathsf{D} = [-20\ 20] \times [-40\ 40]$, $\eta = 0.2$, $\mu = 25$, and $\tau = 300$ seconds. Note that we selected the disturbance set D in such a way that it contains the disturbances signal ω in (II.6) with a large confidence level. We have used the pole placement method [2] in order to design a state observer by choosing the poles of the observer to be -100, -101, -102, and -103. Note that using the proposed observer, the constant ϵ in the definition of $S_{\mathsf{q}}(\Sigma)$ is much smaller than η and is being neglected here.

The resulting cardinality of the state and input sets for $S_{\mathsf{q}}(\Sigma)$ are 10119392 and 12, respectively. The CPU time taken for synthesizing the controller has amounted to 3543 seconds. Figure 3 displays the outputs of the closed loop system (i.e. ζ_1 and ζ_2) stemming from the initial condition $x_0 = [3,\ 0,\ -6.5,\ 0]^T$ while the observer is initialized at $\hat{x}_0 = [-1.5,\ 0,\ -5,\ 0]^T$, the corresponding evolution of the control input signals (i.e. υ_1 and υ_2), as well as the disturbance input ω_2, generated as in (II.6). As can be seen in Figure 3, the outputs of the HVAC system, i.e., $[\zeta_1,\ \zeta_2]^T$, enter the set W in finite time and stay there forever. In the simulation, we considered the disturbance only on the 2nd input and the disturbance on the 1st input to be zero. Although for the sake of constructing $S_{\mathsf{q}}(\Sigma)$, we assumed that the disturbance inputs are piecewise constant of duration τ, in the simulation (c.f. Figure 3) we did not impose such assumption.

Although we have decided for simplicity to deal with the disturbance as a nondeterministic quantity, we can approach this control synthesis problem in the context of stochastic models as in [26], [28], or [17] for discrete-time models. This alternative task is left to future work.

[4]The LTL semantics are defined over the output behaviours of $S_{\mathsf{q}}(\Sigma)$.

IV. Conclusions

We plan to pursue practical generalization, targeting the key aspects for HVAC controllers that have been discussed in the Introduction. For instance, we plan to extend the work by considering the switching between controllers designed for different operating conditions or by synthesizing a controller, in the context of switching stochastic systems [26], which is robust with respect to randomly changing operating conditions. In view of computational complexity issues, we further plan to reduce the number of quantization levels to reduce memory requirements (cf. [27], [30] for potential approaches).

V. Notations

If A is a subset of B we denote by $\imath_A : A \hookrightarrow B$ or simply by $\imath$ the natural inclusion map taking any $a \in A$ to $\imath(a) = a \in B$. The symbols $\mathbb{N}$, $\mathbb{N}_0$, $\mathbb{Z}$, $\mathbb{R}$, $\mathbb{R}^+$, and $\mathbb{R}_0^+$ denote the set of natural, nonnegative integer, integer, real, positive, and nonnegative real numbers, respectively. The symbol $0_{n\times m}$ denotes the zero matrix in $\mathbb{R}^{n\times m}$. Given a vector $x \in \mathbb{R}^n$, we denote by x_i the i–th element of x, and by $\|x\|$ the infinity norm of x, namely $\|x\| = \max\{|x_1|, |x_2|, ..., |x_n|\}$, where $|x_i|$ denotes the absolute value of x_i.

The closed ball centred at $x \in \mathbb{R}^n$ with radius λ is defined by $\mathcal{B}_\lambda(x) = \{y \in \mathbb{R}^n \mid \|x - y\| \leq \lambda\}$. A set $B \subseteq \mathbb{R}^n$ is called a *box* if $B = \prod_{i=1}^n [c_i, d_i]$, where $c_i, d_i \in \mathbb{R}$ with $c_i < d_i$ for each $i \in \{1, \ldots, n\}$. The *span* of a box B is defined as $span(B) = \min\{|d_i - c_i| \mid i = 1, \ldots, n\}$. By defining $[\mathbb{R}^n]_\eta = \{a \in \mathbb{R}^n \mid a_i = k_i\eta, k_i \in \mathbb{Z}, i = 1, \ldots, n\}$, the set $\bigcup_{p\in[\mathbb{R}^n]_\eta} \mathcal{B}_\lambda(p)$ is a countable covering of $\mathbb{R}^n$ for any $\eta \in \mathbb{R}^+$ and $\lambda \geq \eta/2$. For a box $B \subseteq \mathbb{R}^n$ and $\eta \leq span(B)$, define the η-approximation $[B]_\eta = [\mathbb{R}^n]_\eta \cap B$. Note that $[B]_\eta \neq \varnothing$ for any $\eta \leq span(B)$. Geometrically, for any $\eta \in \mathbb{R}^+$ with $\eta \leq span(B)$ and $\lambda \geq \eta$, the collection of sets $\{\mathcal{B}_\lambda(p)\}_{p\in[B]_\eta}$ is a finite covering of B, i.e., $B \subseteq \bigcup_{p\in[B]_\eta} \mathcal{B}_\lambda(p)$. We extend the notions of *span* and of *approximation* to finite unions of boxes as follows. Let $A = \bigcup_{j=1}^M A_j$, where each A_j is a box. Define $span(A) = \min\{span(A_j) \mid j = 1, \ldots, M\}$, and for any $\eta \leq span(A)$, define $[A]_\eta = \bigcup_{j=1}^M [A_j]_\eta$.

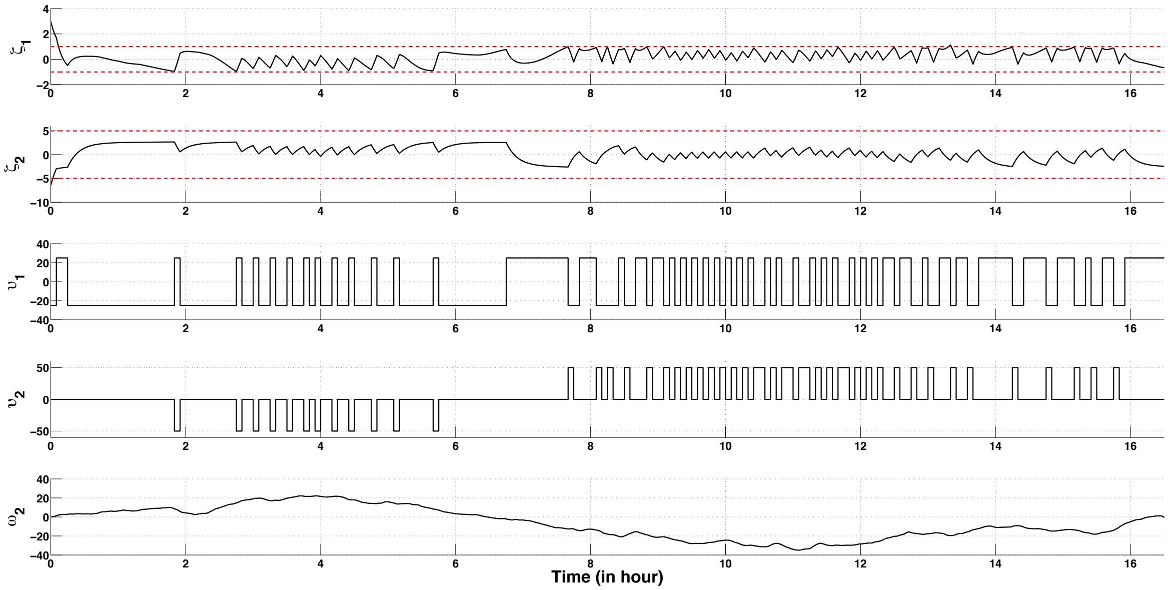

Fig. 3. The outputs of the closed-loop system (top two panels) and the evolution of the input and disturbance signals (bottom three panels). The dashed (red) lines in the top two panels denote the target ranges (c.f. Table II).

REFERENCES

[1] M. Anderson, M. Buehner, P. Young, D. Hittle, C. Anderson, J. Tu, and D. Hodgson. Mimo robust control for hvac systems. *Control Systems Technology, IEEE Transactions on*, 16(3):475–483, 2008.

[2] P. J. Antsaklis and A. N. Michel. *A Linear Systems Primer*. Springer, 2007.

[3] C. Baier and J. P. Katoen. *Principles of model checking*. The MIT Press, April 2008.

[4] S Bengea, A Kelman, Francesco Borrelli, Russell Taylor, and Satish Narayanan. Model predictive control for mid-size commerical building hvac: Implementation, results and energy savings. In *Second International Conference on Building Energy and Environment*, 2012.

[5] V. Bicik, O. Holub, K. Marik M., Sikora, P. Stluka, and R. D'hulst. Platform for coordination of energy generation and consumption in residential neighborhoods. In *Innovative Smart Grid Technologies (ISGT Europe), 2012 3rd IEEE PES International Conference and Exhibition on*, pages 1–7. IEEE, 2012.

[6] A. Kelman F. and Borrelli. Bilinear model predictive control of a hvac system using sequential quadratic programming. In *IFAC World Congress*, 2011.

[7] R. S. Fargus and C. Chapman. A commercial pi-neural controller for the control of building services plant. In *Control'98. UKACC International Conference on (Conf. Publ. No. 455)*, pages 1688–1693. IET, 1998.

[8] ASHRAE Handbook et al. Fundamentals. *American Society of Heating, Refrigerating and Air Conditioning Engineers, Atlanta*, 111, 2001.

[9] Lennox International Inc. Multi-stage air volume. Electronically available at: http://www.lennoxcommercial.com/landing/msav.asp.

[10] Trane Inc. Intellipak i commercial packaged rooftop units. Electronically available at: http://www.trane.com/commercial.

[11] J. Laustsen. Energy efficiency requirements in building codes, energy efficiency policies for new buildings. *International Energy Agency (IEA)*, pages 477–488, 2008.

[12] Trend Control Systems Ltd. Trend makes energy saving a commercial venture. Electronically available at: https://www.trendcontrols.com/en-GB/marketsectors/commercial/Pages/default.aspx.

[13] O. Maler, A. Pnueli, and J. Sifakis. On the synthesis of discrete controllers for timed systems. In E. W. Mayr and C. Puech, editors, *Symposium on Theoretical Aspects of Computer Science*, volume 900 of *LNCS*, pages 229–242. Springer-Verlag, 1995.

[14] K. Mařík, J. Rojíček, P. Stluka, and J. Vass. Advanced hvac control: Theory vs. reality. In *Preprints of the 18th IVAC World Congress, Milano, Italy*, pages 3108–3113, 2011.

[15] M. Mazo Jr., A. Davitian, and P. Tabuada. PESSOA: A tool for embedded control software synthesis. In T. Touili, B. Cook, and P. Jackson, editors, *Computer Aided Verification (CAV)*, volume 6174 of *LNCS*, pages 566–569. Springer-Verlag, July 2010.

[16] L. Pérez-Lombard, J. Ortiz, and C. Pout. A review on buildings energy consumption information. *Energy and buildings*, 40(3):394–398, 2008.

[17] S. Esmaeil Zadeh Soudjani and A. Abate. Adaptive and sequential gridding for abstraction and verification of stochastic processes. *SIAM Journal on Applied Dynamical Systems*, 12(2):921–956, 2012.

[18] S. Esmaeil Zadeh Soudjani and A. Abate. Aggregation and control of populations of thermostatically controlled loads by formal abstractions. *IEEE Transactions on Control Systems Technology*, 23(3):975–990, 2015.

[19] S. Esmaeil Zadeh Soudjani, S. Gerwinn, C. Ellen, M. Fraenzle, and A. Abate. Formal synthesis and validation of inhomogeneous thermostatically controlled loads. In *Quantitative Evaluation of SysTems*, pages 74–89. Springer Verlag, 2014.

[20] ATREA s.r.o. Versatile compact hvac units. Electronically available at: http://www.atrea.cz/en/duplex-500-8000-multi.

[21] P. Stluka. Control for energy-efficient buildings. In *The Impact of Control Technology - 2nd edition*, 2014. Electronically available at: http://ieeecss.org/general/IoCT2-report.

[22] P. Stluka, K. Mařík, and P. Endel. Advanced control solutions for building systems. In *World Congress*, volume 19, pages 606–611, 2014.

[23] C. P. Underwood. Robust control of hvac plant ii: controller design. *Building Services Engineering Research and Technology*, 21(1):63–71, 2000.

[24] C. P. Underwood and F. Yik. *Modelling methods for energy in buildings*. John Wiley & Sons, 2008.

[25] H. H. Asada X. He, Sh. Liu and H. Itoh. Multivariable control of vapor compression systems. *HVAC&R Research*, 4(3):205–230, 1998.

[26] M. Zamani and A. Abate. Approximately bisimilar symbolic models for randomly switched stochastic systems. *Systems and Control Letters*, 69:38–46, July 2014.

[27] M. Zamani, A. Abate, and A. Girard. Symbolic models for stochastic switched systems: A discretization and a discretization-free approach. *Automatica*, 55(5):183–196, 2015.

[28] M. Zamani, P. Mohajerin Esfahani, R. Majumdar, A. Abate, and J. Lygeros. Symbolic control of stochastic systems via approximately bisimilar finite abstractions. *IEEE Transactions on Automatic Control*, 59(12):3135–3150, December 2014.

[29] M. Zamani, G. Pola, M. Mazo Jr., and P. Tabuada. Symbolic models for nonlinear control systems without stability assumptions. *IEEE Transaction on Automatic Control*, 57(7):1804–1809, July 2012.

[30] M. Zamani, I. Tkachev, and A. Abate. Bisimilar symbolic models for stochastic control systems without state-space discretization. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control*, pages 41–50. ACM New York, NY, USA, April 2014.