Implementation of a Scenario-based MPC for HVAC Systems: an Experimental Case Study ⋆

Alessandra Parisio * **Damiano Varagnolo** * **Marco Molinari** * **Giorgio Pattarello** * **Luca Fabietti** ** **Karl H. Johansson** *

* *School of Electrical Engineering, Royal Institute of Technology, Osquldas väg 10, Stockholm, Sweden*
** *Department of Information Engineering, University of Padova, via Gradenigo 6/b, Italy*

Abstract: Heating, Ventilation and Air Conditioning (HVAC) systems play a fundamental role in maintaining acceptable thermal comfort and air quality levels. Model Predictive Control (MPC) techniques are known to bring significant energy savings potential. Developing effective MPC-based control strategies for HVAC systems is nontrivial since buildings dynamics are nonlinear and influenced by various uncertainties. This complicates the use of MPC techniques in practice. We propose to address this issue by designing a stochastic MPC strategy that dynamically learns the statistics of the building occupancy patterns and weather conditions. The main advantage of this method is the absence of a-priori assumptions on the distributions of the uncertain variables, and that it can be applied to any type of building. We investigate the practical implementation of the proposed MPC controller on a student laboratory, showing its effectiveness and computational tractability.

Keywords: Control applications, Implementation, Model-based and predictive control, Probabilistic models, Control-oriented models, Stochastic control

1. INTRODUCTION

Heating, cooling and air conditioning is a necessity in buildings (commercial, residential, and industrial), which account for a major share of the global energy consumptions. Reports indicate that Heating, Ventilation and Air Conditioning (HVAC) systems in developed countries contribute for approximately one fifth of the total national energy usages (European Commission, 2008).

It is generally accepted that buildings frequently use more energy than expected or desired, and that often HVAC control systems do not operate properly. In brief, current practice shows its limits, with potential energy savings achievable by using systematic building management being estimated from 5% to 30% of the total consumptions (Costa et al., 2013; Chua et al., 2013).

These figures indicate a tremendous potentials of improvements. It is not surprising thus that several academic and industrial research groups are actively working on achieving these improvements.

⋆ The research leading to these results has received funding from the European Union Seventh Framework Programme [FP7/2007-2013] under grant agreement n°257462 HYCON2 Network of excellence, the European Institute of Technology (EIT) Information and Communication Technology (ICT) Labs, the Swedish Energy Agency, the Swedish Governmental Agency for Innovation Systems (VINNOVA), the Swedish Research Council and the Knut and Alice Wallenberg Foundation.
Corresponding author: Alessandra Parisio, parisio@kth.se

1.1 Literature Review

The current trend is to improve HVAC control systems performance by using predictive strategies, like in Erickson et al. (2009); Goyal et al. (2012); Gwerder and Toedtli (2005); Salsbury et al. (2012); Hua and Karavab (2014).

This tendency is supported not only by simulations (Treado and Chen, 2013; Wallace et al., 2012; Fadzli Haniff et al., 2013), but also by some experimental results on real buildings (Sturzenegger et al., 2013; Široký et al., 2011; Parisio et al., 2013b). Model Predictive Controls (MPCs) may truly yield better comfort levels and energy use performance than current practices do. Energy savings depend on the specific building and its various factors such as insulation, weather, building occupancy patterns, etc.

In any case, contributions of MPC schemes to energy savings are so relevant that this technology is expected to become a common solution for smart buildings in a few years (Aswani et al., 2012).

In particular, successful implementations will be likely based on stochastic MPC schemes with probabilistic constraints: indeed, indoor air conditions are intrinsically affected by stochastic disturbances, such as unpredictable occupancy patterns and external temperature levels. Current standards thus explicitly accounts for the possibility of comfort violations, stating that the probability of these violations should do not exceed certain levels (BSI, 2008).

There is already a vast literature on stochastic MPC schemes for HVAC control. For example, Mady et al. (2011) considers stochastic occupancy models, while Ma and Borrelli (2012); Ma et al. (2012a,b) propose stochastic predictive regulators with weather and load disturbances modeled as Gaussian processes, and nonlinear programs for the control of indoor temperatures, which are solved with tailored sequential Quadratic Programs (QPs). Oldewurtel et al. (2012) instead considers stochastic weather predictions (but deterministic predictions of the internal gains) and computes control actions by solving linearized non-convex problems and disturbance feedbacks.

Noticeably, all the previously mentioned approaches restrict disturbances to having Gaussian distribution, assumption that simplifies computations and make the problems solvable. Instead, we proposed a scenario-based tractable approximation of the chance constrained MPC problem, where the scenarios are i.i.d. samples extracted from general probability distributions, thus not restricted to be Gaussian (Parisio et al., 2013a,b). Another scenario-based approach has been proposed by Zhang et al. (2013). Here authors propose an iterative bilinearization of the building model around nominal trajectories and sample occupancy scenarios from a set of measurement data collected in eight single offices equipped with motion sensors. The numerical simulations performed in this work suggest that scenarios-based techniques outperform other predictive methods.

We eventually mention that a promising research direction is to deal with the problem of robustify MPC with respect to uncertainties on the building model, due, e.g., to imperfect predictions of internal and external heat gains (Maasoumy and Sangiovanni-Vincentelli, 2012). Here authors present a model predictive control strategy that is robust against additive uncertainty, introduced as imperfect weather and occupancy predictions.

1.2 Statement of contributions

Implementation of prototypes of stochastic MPC HVAC schemes in real buildings has shown to be possible. In this study we primarily focus on investigating the technology readiness level by adopting technological improvements suggested by the previously gained experience. We then describe new methodological and practical implementation details, and report a detailed analysis of the experimental results.

With respect to the state-of-the-art literature, we: *i)* propose a novel building model, which better captures the building dynamics while maintaining linearity assumptions; *ii)* develop and implement on a real testbed (the KTH HVAC testbed) an advanced control scheme that continuously adapt the operation of the HVAC system to unknown disturbances while guaranteeing occupants comfort and wellbeing. More precisely, the new model accounts for minimal ventilation levels and more precise actuators dynamics. Actuators are also now controlled with more sophisticated and better performing control laws.

We then compare and analyze the energy usage of 3 control schemes applied to the KTH HVAC testbed. The first controller is the current practice in our building. The second is a deterministic MPC disregarding information on the uncertainties of the disturbances. The third controller is instead our novel Scenario-based Model Predictive Control (SMPC) scheme. Results show that, as expected, the SMPC scheme leads to a more robust and potentially energy efficient behavior of the system.

1.3 Structure of the manuscript

Section 2 presents the novel building model and related HVAC MPC scheme. Section 3 then describes our experimental campaign, analyzes the energy usage of the various controllers exploited during the data collection phase, and provides remarks on the degree of precision of the numerical results. Section 4 ends the manuscript with a summary of our conclusions and with indications of the next steps.

2. SCENARIO-BASED MPC FOR HVAC SYSTEMS

In this section we first describe the model of the building (Section (2.1)), and then we outline the general structure of the Scenario-based Model Predictive Control (SMPC) control scheme (Section (2.2)).

The inputs of a Model Predictive Control (MPC) scheme for building climate control are, at every time step, *i)* occupancy levels, *ii)* weather conditions, and *iii)* measurements of the current state of the building. The output is instead a heating, cooling and ventilation plan for the next N hours, where N is the prediction horizon. Conforming with the MPC paradigm, only the first step of this control plan is applied to the Heating, Ventilation and Air Conditioning (HVAC) system. After that, the whole procedure is repeated. This introduces feedback into the system, since the control action is a function of the current state and currently acting disturbances. In our case the computed outputs are, at every time step k, *i)* a mass air flow rate $\dot{m}_{\mathrm{venting}}(k)$, *ii)* a supply air temperature $T_{\mathrm{sa}}(k)$, and *iii)* a radiators mean radiant temperature T_{mr}.

We notice that, since the overall building energy usage is commonly computed as the sum the energy usages of the single thermal zones Gwerder and Toedtli (2005), here we focus on the control of a single thermal zone (or room).

2.1 Modeling

To improve the computational tractability of the overall control problem, we take advantage from the independence of the CO_2 concentration dynamics from the thermal ones, which allow us to address two separated subproblems: *i)* the CO_2-SMPC problem, which aims at minimizing energy use while keeping CO_2 levels in given comfort bounds; *ii)* the T-SMPC problem, controlling instead the indoor temperature.

Here we describe the two separated models for the dynamics under consideration.

Model for the CO_2 concentration dynamics The model is derived from a CO_2 balance equation accounting for the fresh air from the ventilation system and the amount of CO_2 generated per occupant. The state of the model and its output, indicated respectively with x_{CO_2} and y_{CO_2}, are set to be equal to $\Delta\mathrm{CO}_2$, the nonnegative difference between the CO_2 concentration in the room and the inlet air CO_2 concentration (the latter assumed equal to outdoor CO_2 concentration levels).

The model disturbance $w_{\mathrm{CO}_2}(k)$ represents the number of occupants, while the control input is the rate of the air flow coming from the ventilation system, which is denoted by $\dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2}$. This input allows to control the heat flow due to the ventilation system, indicated with Q_{venting}.

The reduction in the indoor CO_2 concentration levels induced by $\dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2}$ is modeled with the bilinear term $\dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2} \cdot x_{\mathrm{CO}_2}$. Since linear problems can be solved more efficiently than nonlinear ones, we derive an equivalent linear model of the CO_2 concentration dynamics by introducing the auxiliary input $u_{\mathrm{CO}_2} := \dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2} \cdot x_{\mathrm{CO}_2}$, which then hides the bilinear term defined above. To meet the physical bounds on the original control input $\dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2}$, u_{CO_2} has to satisfy

$$\dot{m}_{\mathrm{venting}}^{\min} \cdot x_{\mathrm{CO}_2}(k) \leq u_{\mathrm{CO}_2}(k) \leq \dot{m}_{\mathrm{venting}}^{\max} \cdot x_{\mathrm{CO}_2}(k). \quad (1)$$

Then, we can then easily derive $\dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2}(k)$ by inverting the definition of $u_{\mathrm{CO}_2}(k)$.

We assume bounds on the input $u_{\mathrm{CO}_2}(k)$ of the form $u_{\mathrm{CO}_2}^{\min} \leq u_{\mathrm{CO}_2}(k) \leq u_{\mathrm{CO}_2}^{\max}$. These bounds can be expressed as polytopic constraints $F u_{\mathrm{CO}_2}(k) \leq f$. We further define comfort constraints on the indoor CO_2 concentration as $0 \leq y_{\mathrm{CO}_2}(k) \leq y_{\mathrm{CO}_2}^{\max}$. Considering that $x_{\mathrm{CO}_2} = y_{\mathrm{CO}_2}$, comfort constraints and constraints (1) can be written in a compact form as mixed constraints on the input and on the output, $V_y y_{\mathrm{CO}_2}(k) + V_u u_{\mathrm{CO}_2}(k) \leq v$. We refer the reader to Parisio et al. (2013b) for details on the construction of the constraints matrices.

With the control input u_{CO_2}, the CO_2 concentration dynamics can eventually be described by the discrete-time Linear Time Invariant (LTI) system

$$\begin{aligned} x_{\mathrm{CO}_2}(k+1) &= a x_{\mathrm{CO}_2}(k) + b u_{\mathrm{CO}_2}(k) + e w_{\mathrm{CO}_2}(k) \\ y_{\mathrm{CO}_2}(k) &= x_{\mathrm{CO}_2}(k). \end{aligned} \quad (2)$$

Model for the thermal dynamics We consider a thermal Resistive-Capacitive (RC) network of first-order systems, where the nodes are the states representing the temperatures of the room, walls, floor and ceiling. Each state is associated to a heat transfer differential equation.

The model disturbances represent the outdoor temperature, radiation, internal gains, heat flows due to occupancy, equipments and lightings. The control inputs are the temperature of the supplied air, T_{sa}, the mean radiant temperature of the radiators, T_{mr}, and the air flow rate $\dot{m}_{\mathrm{venting}}$. (We remind that $\dot{m}_{\mathrm{venting}}$ must be at least equal to $\dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2}$, the latter representing the minimum air flow rate needed to maintain optimal CO_2 levels.) The inputs T_{sa}, T_{mr} and $\dot{m}_{\mathrm{venting}}$ allow to control two different heat flows: *i)* Q_{venting}, representing the contribute due to the ventilation system; *ii)* Q_{heating}, representing the contribute due to the radiators.

We now aim to: *i)* hide the bilinear term of the indoor thermal dynamics $Q_{\mathrm{venting}} = \dot{m}_{\mathrm{venting}} c_{\mathrm{pa}} (T_{\mathrm{sa}} - T_{\mathrm{room}})$, *ii)* model the contribute due to the requirements on the CO_2 concentration levels (i.e., due to the minimal air flow $\dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2}$) and the absolute value of Q_{venting} (which is part of the cost function to be minimized).

To achieve these aims, we model the two heat flows as

$$\begin{aligned} Q_{\mathrm{venting}} &= \dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2} c_{\mathrm{pa}} \big(\Delta T_h - \Delta T_c\big) + c_{\mathrm{pa}} \big(\Delta u_h - \Delta u_c\big) \\ Q_{\mathrm{heating}} &= A_{\mathrm{rad}} h_{\mathrm{rad}} \Delta T_{\mathrm{h,rad}} \end{aligned}$$

where c_{pa} is the specific heat of the dry air, A_{rad} is the emission area of the radiators, h_{rad} is the heat transfer coefficient of the radiators, and the nonnegative variables ΔT_h, ΔT_c, Δu_h and Δu_c are s.t.

$$\begin{aligned} \Delta T_h - \Delta T_c &= T_{\mathrm{sa}} - T_{\mathrm{room}} \\ \Delta T_h + \Delta T_c &= \big|T_{\mathrm{sa}} - T_{\mathrm{room}}\big| \\ \Delta u_h - \Delta u_c &= \Delta \dot{m}_{\mathrm{venting}} \big(T_{\mathrm{sa}} - T_{\mathrm{room}}\big) \\ \Delta u_h + \Delta u_c &= \Delta \dot{m}_{\mathrm{venting}} \big|T_{\mathrm{sa}} - T_{\mathrm{room}}\big| \end{aligned}$$

with $\Delta \dot{m}_{\mathrm{venting}} := \dot{m}_{\mathrm{venting}} - \dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2}$ the additional air flow rate required for guaranteeing the thermal comfort, and T_{room} the indoor temperature.

We then represent physical bounds on the original control inputs as

$$T_{\mathrm{sa}}^{\min} - T_{\mathrm{room}}(k) \leq \Delta T_h(k) - \Delta T_c(k) \leq T_{\mathrm{sa}}^{\max} - T_{\mathrm{room}}(k) \quad (3)$$

$$\big|\Delta u_h(k) - \Delta u_c(k)\big| \leq \Delta \dot{m}_{\mathrm{venting}}^{\max}(k) \big|\Delta T_h(k) - \Delta T_c(k)\big| \quad (4)$$

where $\Delta \dot{m}_{\mathrm{venting}}^{\max}(k) := \dot{m}_{\mathrm{venting}}^{\max} - \dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2}(k)$.

With the newly introduced variables, the dynamics of the indoor temperature can be modeled with the discrete-time Linear Time Invariant (LTI) system

$$\begin{aligned} x_{\mathrm{T}}(k+1) &= A_{\mathrm{T}} x_{\mathrm{T}}(k) + B_{\mathrm{T}}(k) u_{\mathrm{T}}(k) + E_{\mathrm{T}} w_{\mathrm{T}}(k) \\ y_{\mathrm{T}}(k) &= C_{\mathrm{T}} x_{\mathrm{T}}(k), \end{aligned} \quad (5)$$

where the state $x_{\mathrm{T}}(k)$ contains the temperatures of the room and of the inner and outer parts of the walls, $u_{\mathrm{T}}(k) := \Big[\Delta T_h(k), \Delta T_c(k), \Delta u_h(k), \Delta u_c(k), \Delta T_{h,\mathrm{rad}}(k)\Big]$ is the input vector, and $w_{\mathrm{T}}(k)$ is the vector of random disturbances (outdoor temperature, solar radiation and internal heat gains). The output $y_{\mathrm{T}}(k)$ is the indoor temperature at time k. We notice that the input matrix $B_{\mathrm{T}}(k)$ is time varying since it depends on $\dot{m}_{\mathrm{venting}}^{\mathrm{CO}_2}(k)$.

Compared to our previous contributions (Parisio et al., 2013a,b), the building model now encompasses a more detailed solar radiation model. Furthermore, the temperature variation in adjacent rooms has been estimated by means of a sinusoidal dependence in time, which proved to be in sufficiently good accordance with measured data.

By following the same procedure outlined in sub-section 2.1, we impose hard constraints on the inputs and comfort constraints on the indoor temperature. Hence, hard constraints on inputs and constraints (4) can be written in compact form as polytopic constraints on inputs, $F u_{\mathrm{T}}(k) \leq f$. Comfort constraints on the output and constraints (3) can be written in a compact form as mixed constraints on the input and on the output, $V_y y_{\mathrm{T}}(k) + V_u u_{\mathrm{T}}(k) \leq v$.

We eventually notice that, once $\dot{m}^{\mathrm{CO_2}}_{\mathrm{venting}}(k)$ and $u_{\mathrm{T}}(k)$ have been computed, the original control variables $T_{\mathrm{sa}}(k)$, $T_{\mathrm{mr}}(k)$ and $\dot{m}_{\mathrm{venting}}(k)$ can be easily computed by simple inversion formulas.

2.2 Scenario-based MPC

As suggested in the modeling section, we decouple the synthesis problem in two separated parts and formulate two problems: the CO_2-SMPC problem, which considers model (2), and the T-SMPC problem, which includes model (5).

We also remark that, since the requirements on CO_2 concentrations have priority over the thermal comfort ones, the solution computed by the CO_2-SMPC is considered by the T-SMPC as a lower bound on the air flow rate.

We thus consider an MPC problem for the control of discrete-time LTI systems of the form

$$\begin{aligned} x(k+1) &= Ax(k) + B(k)u(k) + Ew(k) \\ y(k) &= Cx(k), \end{aligned} \tag{6}$$

where $x(k) \in \mathcal{R}^n$ is the state, $u(k) \in \mathcal{R}^m$ is the control input, $w(k) \in \mathcal{R}^r$ is the stochastic disturbance and $y(k) \in \mathcal{R}^p$ is the output. Indeed (6) represents either (2) or (5), depending on the controller under consideration (CO_2-SMPC or T-SMPC).

Consider then a prediction horizon N and define

$$\begin{aligned} \boldsymbol{x} &:= \left[x(1)^{\mathrm{T}}, \ldots, x(N)^{\mathrm{T}}\right]^{\mathrm{T}}, \\ \boldsymbol{u} &:= \left[u(0)^{\mathrm{T}}, \ldots, u(N-1)^{\mathrm{T}}\right]^{\mathrm{T}}, \\ \boldsymbol{y} &:= \left[y(0)^{\mathrm{T}}, \ldots, y(N-1)^{\mathrm{T}}\right]^{\mathrm{T}}, \\ \boldsymbol{w} &:= \left[w(0)^{\mathrm{T}}, \ldots, w(N-1)^{\mathrm{T}}\right]^{\mathrm{T}}, \end{aligned}$$

where $x(k+1) = Ax(k) + Bu(k) + Ew(k)$ denotes the predictions of the state after k time instants into the future. Defining the prediction dynamics matrices $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{E}$ and $\boldsymbol{C}$ s.t.

$$\begin{aligned} \boldsymbol{x} &= \boldsymbol{A}x(0) + \boldsymbol{B}\boldsymbol{u} + \boldsymbol{E}\boldsymbol{w} \\ \boldsymbol{y} &= \boldsymbol{C}\boldsymbol{x}, \end{aligned} \tag{7}$$

with $x(0)$ the current measured value of the state, we can express the output as a function of the initial state $x(0)$, i.e., as

$$\boldsymbol{y} = \boldsymbol{C}_A x(0) + \boldsymbol{C}_B \boldsymbol{u} + \boldsymbol{C}_E \boldsymbol{w}. \tag{8}$$

The linear constraints on the inputs and outputs over the prediction horizon can instead be generally written as

$$\begin{aligned} \boldsymbol{V}_y \boldsymbol{y} + \boldsymbol{V}_u \boldsymbol{u} &\leq \boldsymbol{v} \\ \boldsymbol{F}\boldsymbol{u} &\leq \boldsymbol{f}, \end{aligned} \tag{9}$$

where $\boldsymbol{F} \in \mathcal{R}^{q \times mN}$, $\boldsymbol{f} \in \mathcal{R}^q$, $\boldsymbol{V}_y \in \mathcal{R}^{r \times pN}$, $\boldsymbol{V}_u \in \mathcal{R}^{r \times mN}$ and $\boldsymbol{v} \in \mathcal{R}^r$.

By replacing (8) in (9), we can write the constraints on the outputs as $\boldsymbol{G}_u \boldsymbol{u} + \boldsymbol{G}_w \boldsymbol{w} \leq \boldsymbol{g}$, where $\boldsymbol{G}_u$, $\boldsymbol{G}_w$ and $\boldsymbol{g}$ are matrices of appropriate dimensions.

MPCs can then be formulated so that it can simultaneously incorporate weather and occupancy forecasts and their uncertainties by means of *chance-constrained* formulations. It is indeed possible to assume the possibility of violating the comfort bounds on the indoor temperature and CO_2 levels with a predefined probability, i.e., formulate output constraints as

$$\mathbb{P}\left[\boldsymbol{G}_u \boldsymbol{u} + \boldsymbol{G}_w \boldsymbol{w} \leq \boldsymbol{g}\right] \geq 1 - \alpha.$$

with $\alpha \in [0, 1]$ being the violation probability level. In these formulations α represents a tradeoff between performance and constraint satisfaction.

The cost function represents the energy use over the whole prediction horizon. Denoting by $\boldsymbol{c}^{\mathrm{T}} \boldsymbol{u} \Delta k$ with $\boldsymbol{c} \in \mathbb{R}^{mN}$ the cost vector and Δk the sampling period, the control problem can be formally stated as

Problem 1. (Chance Constrained MPC for HVAC Control).

$$\begin{aligned} \min_{\boldsymbol{y}} \ & \boldsymbol{c}^{\mathrm{T}} \boldsymbol{u} \Delta k \\ \text{s.t.} \ & \mathbb{P}\left[\boldsymbol{G}_u \boldsymbol{u} + \boldsymbol{G}_w \boldsymbol{w} \leq \boldsymbol{g}\right] \geq 1 - \alpha \\ & \boldsymbol{F}\boldsymbol{u} \leq \boldsymbol{f}, \end{aligned}$$

with $1 - \alpha$ the desired probability level for constraint satisfaction.

Chance constrained problems like 1 are generally intractable unless the uncertainties follow specific distributions, e.g., Gaussian or log-concave. In these cases it is possible to obtain equivalent convex (and thus computationally efficient) reformulations, as in Kall and Mayer (2005).

However, as described later, Gaussian assumptions are rather restrictive. To overcome this limitation but still obtain a solvable MPC problem we propose to apply *randomized* approaches Calafiore (2010), which do not require the specification of particular probability distributions for the uncertainties but only the capability of randomly extracting from them.

The approach is as follows: *i)* let $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_S$ be a set of S i.i.d. disturbances samples (called *scenarios*), $\boldsymbol{w}_i := \left[w_i^{\mathrm{T}}(0), \ldots, w_i^{\mathrm{T}}(N-1)\right]^{\mathrm{T}}$, $i = 1, \ldots, S$. Then the chance constraints in Problem (1) can be replaced with the deterministic constraints

$$\boldsymbol{G}_u \boldsymbol{u} + \boldsymbol{G}_w \boldsymbol{w}_i - \boldsymbol{g} \leq \boldsymbol{0}, \quad i = 1, \ldots, S.$$

Since most of the constraints in (10) are redundant, the only constraint that is required to be satisfied is

$$\boldsymbol{G}_u \boldsymbol{u} \leq \boldsymbol{g} - \max_{i=1,\ldots,S} \boldsymbol{G}_w \boldsymbol{w}_i$$

(where the max applies element-wise to $\boldsymbol{G}_w \boldsymbol{w}_i$); *ii)* soften the constraints in (10) by introducing the slack variables $\epsilon(k) \in \mathbb{R}^p$ at each time step k, and eventually approximate Problem 1 with

Problem 2. (SMPC for HVAC Control).

$$\begin{aligned} \min_{\boldsymbol{u}} \ & \boldsymbol{c}^{\mathrm{T}} \boldsymbol{u} \Delta k + \rho \boldsymbol{1}^{\mathrm{T}} \boldsymbol{\epsilon} \\ \text{s.t.} \ & \boldsymbol{G}_u \boldsymbol{u} \leq \boldsymbol{g} + \boldsymbol{\epsilon} - \max_{i=1,\ldots S} \boldsymbol{G}_w \boldsymbol{w}_i \\ & \boldsymbol{F}\boldsymbol{u} \leq \boldsymbol{f}, \end{aligned} \tag{10}$$

where $\boldsymbol{\epsilon}$ is the vector containing all the slack variables, ρ is the weight on the slack variables, and $\boldsymbol{1}$ is a matrix of ones with appropriate dimensions.

We notice two important remarks:

- (how to choose the number of scenarios S) letting $d = mN$ be the number of decision variables, S can be chosen based on the sufficient condition
$$S \geq \frac{2}{\alpha}\left(\ln\left(\frac{1}{\beta}\right) + d\right), \tag{11}$$
that guarantees that considering constraints (10) will lead to a feasible solution for Problem 2 with a confidence level $(1-\beta) \in (0,1)$ with β an user-defined parameter (Calafiore, 2010; Calafiore and Campi, 2008). Our experience nonetheless indicates that (11) may be overly pessimistic. E.g., we ran numerical simulations with $\alpha = 0.05$ and $\beta = 0.001$ and computed the empirical probability of constraint violation over 2400 different i.i.d. instances of the random convex problem (2). Applying condition (11), we set $M = 3157$ and empirically reported a constraints violations probability of 0.0044. Halving the indication given by (11) ($M = 1579$) instead led to an empirical probability of constraint violations of 0.042, much closer to the confidence level required initially.
- (meaning of the slack variables $\epsilon(k)$'s) the $\epsilon(k)$'s tune the number of possible constraint violations and guarantee that the problem with sampled constraints is always feasible. If the optimal solution can be obtained without violations of the softened constraints, the slack variables will be set to zero. The designer can thus considerably penalize constraint violations by assigning to the weighting factor a value that is orders of magnitude greater than the other coefficients parameters.

For algorithms for the generation of the scenarios we send back the interested reader to Parisio et al. (2013a,b).

3. EXPERIMENTAL CASE STUDY

3.1 Description of the experimental setup

We consider a laboratory of approximatively $80m^2$ in the ground floor of the Q-building of the KTH Royal Institute of Technology campus in Stockholm. The room has a concrete heavyweight structure with limited glass surface and one external wall, facing South-East, which is partially shaded by a parking lot. As summarized in Figure 1, its HVAC system is composed mainly of two parts: the ventilation system, supplying fresh air, and a radiator heating system.

The air in the ventilation system is pushed from a central fan (not controllable by us) that is active only between 8:00 and 16:00 during working days. Thus no ventilation control action can be carried when this fan is off, and as a consequence we only report tests performed when the central fan was running. A part from this, the ventilation system works as follows: the balanced ventilation system pre-conditions fresh air from outside, distributing it at a temperature of about 20℃. Part of this generated air flow is then conveyed directly into the room, while part can be further cooled by a cooling coil. Summarizing, the controllable actuators of the ventilation system are 3: two dampers that regulate the opening of the inflow and outflow ducts, and a valve that regulates the temperature

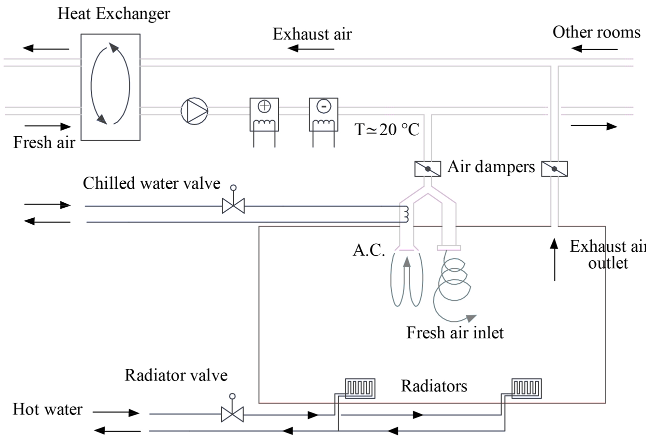

Fig. 1. Scheme of the HVAC system of the testbed.

of the air chilling circuit. When the central fan is on, a minimum level of the air flow rate is guaranteed in any case.

We remark that the ventilation / cooling system is underdimensioned: in case of extreme occupancy levels (e.g., 25 people) and relatively moderately high external temperature (e.g., above 0 ℃), full actuation is not sufficient to maintain internal temperature / CO_2 levels inside the respective comfort bounds, as shown in some of the investigated experimental cases.

The heating system is instead composed by common radiators. The flowing hot water is provided by a district heating, that autonomously decide the temperature of the fluid considering the external temperature conditions. The unique actuator that can be controlled in our testbed is thus the valve regulating the flow of the heating fluid.

Figure 2 depicts the architecture of the implemented control system: the indoor temperature and CO_2 concentration are controlled through the ventilation system and radiators, which are actuated using low-level PI controllers. The set-points for the low-level controllers are computed by our SMPC at each time instant, based on new measurements and updated information about weather and occupancy patterns.

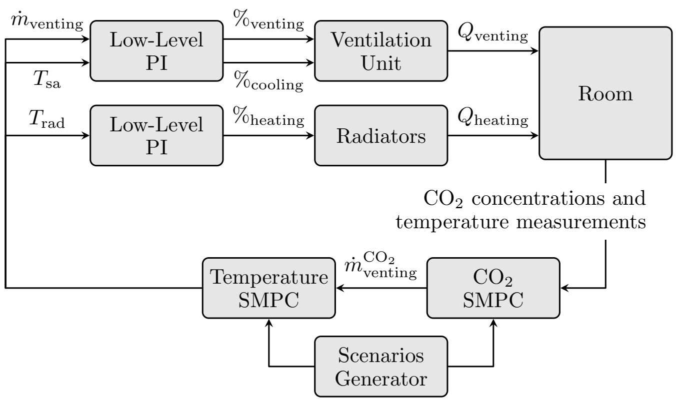

Fig. 2. Architecture of the control system implemented on the testbed.

The scenario-based controller is synthesized according to Algorithm 1.

Algorithm 1 Control Synthesis

for $k = 1, 2, \dots$ **do**
  set $x_{\mathrm{CO_2}}(0) = x_{\mathrm{CO_2}}(k)$ and $x_{\mathrm{T}}(0) = x_{\mathrm{T}}(k)$
  compute the minimum number of scenarios for the CO_2-SMPC problem, $S_{\mathrm{CO_2}}$, and for the T-SMPC problem, S_{T}
  extract $S_{\mathrm{CO_2}}$ occupancy scenarios (for the CO2 control problem) and S_{T} weather and occupancy scenarios (for the temperature control problem) over the prediction horizon N
  solve the CO_2-SMPC problem and compute the sequence $\left\{\dot{m}^{\mathrm{CO_2}}_{\mathrm{venting}}(0), \dots, \dot{m}^{\mathrm{CO_2}}_{\mathrm{venting}}(N-1)\right\}$
  solve the T-SMPC problem and compute $\big(\dot{m}_{\mathrm{venting}}(0), T_{\mathrm{sa}}(0), T_{\mathrm{mr}}(0)\big)$
  set the setpoints of the low-level PI controllers the values computed at the previous step, compute the actuation commands and actuate
end for

3.2 Model Validation

Figures 3 and 4 reports graphical validations of the CO_2 and temperature models (2) and (5) against data from the testbed collected during July 2013. We notice that the models accurately capture the dynamics of the systems in consideration, and that they constitute an improvement w.r.t. the models considered in Parisio et al. (2013b).

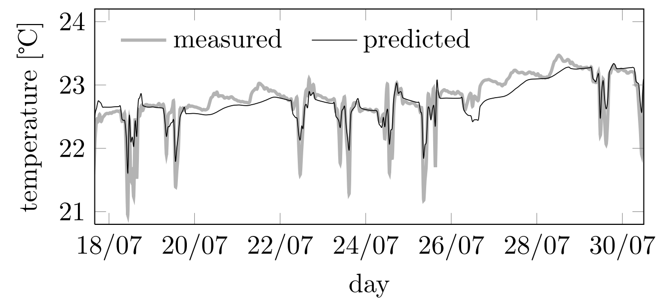

Fig. 3. Validation of the thermal model using the measured temperatures collected from the testbed.

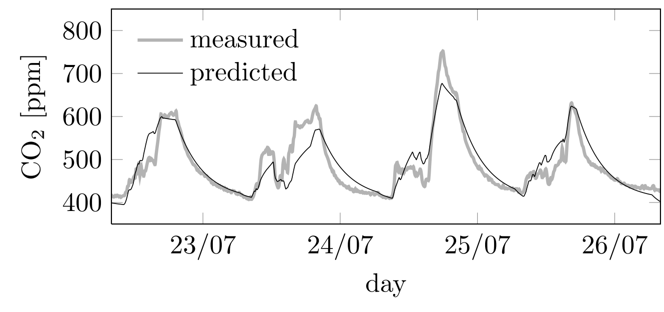

Fig. 4. Validation of the CO_2 concentration model using the measured concentrations collected from the testbed.

3.3 Evaluation of Experimental Results

We consider two performance indexes, *i)* total energy usage, and *ii)* levels of violations of the comfort bounds, respectively calculated as

$$E_{\mathrm{tot}} = c_{\mathrm{pa}} \sum_{k=0}^{N-1} \dot{m}_{\mathrm{venting}}(k) \Big| T_{\mathrm{sa}}(k) - T_{\mathrm{room}}(k) \Big| \Delta k \quad [\mathrm{kWh}] \, ,$$

$$C_h = \sum_{k \text{ s.t. } T_{\mathrm{room}}(k) > T_{\mathrm{UB}}} \Big(T_{\mathrm{room}}(k) - T_{\mathrm{UB}} \Big) \Delta k \quad [°\mathrm{C\ h}] \, .$$

T_{UB} and Δk above are respectively the upper bound on the internal temperature comfort level and the time between two samples.

We compare three controllers: *i)* the current practice, a simple control logic with PI control loops and switching logic, indicated by the acronym "AHC" (from Akademiska Hus, the company managing the building of the testbed). *ii)* a Deterministic Model Predictive Control (DMPC) neglecting information on the uncertainties in the forecasts, and computing the control inputs by solving Problem (1) with deterministic constraints obtained by replacing the unknown disturbances with their forecasts. *iii)* our SMPC.

The controllers are tested between November 3 and 19 2013 for 6 hours each day, from 9:30 to 15:30. The sampling time for the MPC-based controllers is 10 minutes, while the predictions horizon for the weather, occupancy and solar radiance processes is 8 hours. The comfort range of the indoor temperature is $[20, 22]$ °C.

Figure 6 shows experimental results for high-occupancy and low-occupancy cases. Despite the difference in time, the weather conditions during the experiments are similar, as shown in Figure 6, while the occupancy patterns varied during the experimental period. We devise two different occupancy profiles: high and low.

Results for high-occupancy tests (November 13 for AHC, November 7, after 12:30, for DMPC and November 11 for SMPC) Remarkably, for the SMPC case, the outdoor temperature is lower but its effect on the controllers' performance is negligible since the occupancy is the dominating disturbance.

The upper bound on the indoor temperature is violated in all the cases due to the ventilation system being underdimensioned. This leads to quite similar profiles on the indoor CO_2 concentration and temperature for the three controllers during the high occupancy hours. Notice that the SMPC controller has to handle a higher occupancy than the other two controllers. We further point out that, for the AHC controller the supply air temperature exhibits a peak at 13:00 due to the change in the occupancy pattern, while, for the MPC-based controllers, the profile of the supply air temperature from the ventilation system is significantly smoother and does not increase too much. This behavior difference is an example of the added value of the forecasts: the AHC controller does not have knowledge of the upcoming occupancy pattern and decides to turn the ventilation system off at 12:30-13:00, despite the high indoor temperature and the expected number of people.

Results for low-occupancy tests (November 19 for AHC, November 5 for DMPC and November 6 for SMPC) November 19 is a zero-occupancy day and the indoor CO_2

and temperature never violate the comfort ranges, so no control action is necessary.

Considering the two MPC-based controllers' performance in low-occupancy days, we notice that both the disturbances' and CO_2 profiles are really similar, while the control inputs are different. The supply air temperature for DMPC increases by 1 ℃ during the time intervals 12:30-13:30 and 14:30-15:30, and the ventilation system is turned on and off often during the morning (9:30-12:30). The air supply temperature for the scenario-based controller SMPC is smoother and kept lower on average. This behavior is mainly caused by a more stressed precooling effect during the morning (the ventilation system is always kept on from 9:30 to roughly 11:30). This leads to an indoor temperature profile with smaller variations, which is a more favorable behavior in terms of comfort. Further, the indoor temperature for the SMPC controller is kept closer to the lower bound.

The differences in the performance of the two MPC-based controllers emphasize the added value of the incorporated information on the disturbances affecting the system: SMPC computes its control inputs based on a worst-case scenario approach, which leads to a more robust behavior against unknown disturbances.

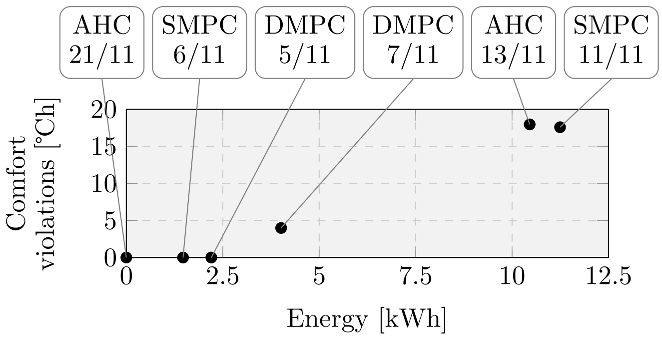

Fig. 5. Comparison of AHC, DMPC and SMPC controllers in terms of total energy usage versus levels of violations of the thermal comfort bounds.

Figure 5 presents the performance indexes of the investigated controllers by plotting their energy usage versus the violations of the thermal comfort bounds. For the high-occupancy cases, DMPC yields to less violations, and uses less thermal energy because of the low-occupancy during the first 3 hours. Comparing the AHC and SMPC controllers, instead, we notice that, despite the higher occupancy, the number of violations are larger for the AHC controller at a cost of a slightly higher energy use.

Looking at the low-occupancy cases, both DMPC and SMPC result in no violations. However the SMPC has a smaller energy use. We point out that one important benefit of SMPC is the possibility of tuning the violation level and then further reduce the energy use.

4. CONCLUSIONS

This paper extends the research line on Scenario-based Model Predictive Control (SMPC) for Heating, Ventilation and Air Conditioning (HVAC) systems started in Parisio et al. (2013a,b) by proposing a novel scenario-based model predictive controller for building climate control.

This control scheme uses weather and occupancy forecasts and takes into account the uncertainty by learning the statistics of the uncertainties on weather and occupancy patterns. With respect to the existing literature, the paper offers three major contributions: *i)* improvements in the modeling of both the building dynamics and its actuators, leading to a novel and more efficient optimization model for Model Predictive Control (MPC) schemes. *ii)* improvements in the practical implementation of the proposed control scheme. *iii)* a detailed analysis on the control performance in terms of energy usage vs. occupants comfort levels on a real building, namely a testbed located in Stockholm, Sweden. More precisely, the document compares the energy usage and the comfort violations of 3 different controllers, cycled for a period of about 3 weeks: the current practice, i.e., the controller normally used by the building's manager; a deterministic MPC; our SMPC.

The experimental results show that there is a promising energy savings potential for SMPC. MPC-based controllers can outperform the current control practice not only in terms of energy usage and comfort levels, but also in terms of more favorable indoor temperature dynamics. In particular, the SMPC leads to temperature variations favorably smaller than the ones obtained with the other control schemes. Also when compared to Deterministic Model Predictive Control (DMPC), SMPC appears to be superior, mainly due to the fact that (unlike DMPC) SMPC is able to directly account for the uncertainty of the weather and occupancy forecast in its control decisions. A further benefit of our SMPC controller is the easy tunability of the tradeoff between energy usage and comfort violations with one tuning parameter describing the level of constraint violations.

We then notice that the proposed SMPC technology is still not completely mature and ready to be massively deployed. Indeed, current implementations require information on the state of the building that up to now are collected using measurement systems usually not present in the majority of the existing buildings (e.g., sensors measuring the temperature of the walls). Thus we devise the necessity of developing advanced estimation schemes that provide indirectly this information. An other important research direction is to extend the control scheme towards networks of thermal zones: the current implementations indeed consider each thermal zone independently and this is inefficient from a optimization problem point of view.

REFERENCES

Aswani, A., Master, N., and Taneja, J. (2012). Energy-efficient building HVAC control using hybrid system LBMPC. In *IFAC Conference on Nonlinear Model Predictive Control*.

BSI (2008). En 15251:2007: Indoor environmental input parameters for design and assessment of energy performance of buildings addressing indoor air quality, thermal environment, lighting and acoustics. Technical report, British Standards Institute.

Calafiore, G. (2010). Random Convex programs. *SIAM Journal on Optimization*, 20(6), 3427–3464.

Calafiore, G. and Campi, M. (2008). The Scenario Approach to Robust Control Design. *IEEE Transactions on Automatic Control*, 51(5), 742–753.

Chua, K., Chou, S., Yang, W., and Yan, J. (2013). Achieving better energy-efficient air conditioning - A review of technologies and strategies. *Applied Energy*, 104, 87–104.

Costa, A., Keane, M.M., Torrens, J.I., and Corry, E. (2013). Building operation and energy performance: Monitoring, analysis and optimisation toolkit. *Applied Energy*, 101, 310–316.

Erickson, V.L., Lin, Y., Kamthe, A., Brahme, R., Surana, A., Cerpa, A.E., Sohn, M.D., and Narayanan, S. (2009). Energy efficient building environment control strategies using real-time occupancy measurements. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings - BuildSys*, 19. ACM Press, New York, New York, USA.

European Commission (2008). Debate Europe-building on the experience of plan D for democracy, dialogue and debate. European Economic and Social Committee and the Committee of the Regions, COM 158/4.

Fadzli Haniff, M., Selamat, H., Yusof, R., Buyamin, S., and Sham Ismail, F. (2013). Review of HVAC scheduling techniques for buildings towards energy-efficient and cost-effective operations. *Renewable and Sustainable Energy Reviews*, 27, 94–103.

Goyal, S., Ingley, H.A., and Barooah, P. (2012). Zone-Level Control Algorithms Based on Occupancy Information for Energy Efficient Buildings. In *American Control Conference*, 3063–3068.

Gwerder, M. and Toedtli, J. (2005). Predictive control for integrated room automation. In *Clima - RHEVA World Congress*.

Hua, J. and Karavab, P. (2014). Model Predictive Control Strategies for Buildings with Mixed-Mode Cooling. *Building and Environment*, 71, 233–244.

Kall, P. and Mayer, J. (2005). *Stochastic Linear Programming: Models, Theory, and Computation*. Springer-Verlag.

Ma, Y. and Borrelli, F. (2012). Fast stochastic predictive control for building temperature regulation. In *American Control Conference*, 3075–3080.

Ma, Y., Borrelli, F., and Hencey, B. (2012a). Model predictive control for the operation of building cooling systems. *IEEE Transactions on Control Systems Technology*, 20(3), 796 – 803.

Ma, Y., Vichik, S., and Borrelli, F. (2012b). Fast stochastic MPC with optimal risk allocation applied to building control systems. In *IEEE Conference on Decision and Control*, 7559–7564. Ieee.

Maasoumy, M. and Sangiovanni-Vincentelli, A. (2012). Optimal control of building HVAC systems in the presence of imperfect predictions. In *Proceedings of the Dynamic System Control Conference*, 1–10.

Mady, A.E.d.D., Provan, G.M., Ryan, C., and Brown, K.N. (2011). Stochastic Model Predictive Controller for the Integration of Building Use and Temperature Regulation. In *AAAI Conference on Artificial Intelligence*, 1371–1376.

Oldewurtel, F., Parisio, A., Jones, C.N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., and Morari, M. (2012). Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, 45(45), 15–27.

Parisio, A., Molinari, M., Varagnolo, D., and Johansson, K.H. (2013a). A scenario-based predictive control approach to Building HVAC managements Systems. In *IEEE Conference on Automation Science and Engineering*.

Parisio, A., Varagnolo, D., Risberg, D., Pattarello, G., Molinari, M., and Johansson, K.H. (2013b). Randomized Model Predictive Control for HVAC Systems. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*.

Salsbury, T., Mhaskar, P., and Qin, S.J. (2012). Predictive control methods to improve energy efficiency and reduce demand in buildings. *Computers & Chemical Engineering*, 51, 77–85.

Sturzenegger, D., Gyalistras, D., Gwerder, M., Sagerschnig, C., Morari, M., and Smith, R.S. (2013). Model Predictive Control of a Swiss office building. In *Clima-RHEVA World Congress*.

Treado, S. and Chen, Y. (2013). Saving Building Energy through Advanced Control Strategies. *Energies*, 6(9), 4769–4785.

Široký, J., Oldewurtel, F., Cigler, J., and Pr´ivara, S. (2011). Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy*, 88(9), 3079–3087.

Wallace, M., McBride, R., Aumi, S., Mhaskar, P., House, J., and Salsbury, T. (2012). Energy efficient model predictive building temperature control. *Chemical Engineering Science*, 69(1), 45–58.

Zhang, X., Schildbach, G., Sturzenegger, D., and Morari, M. (2013). Scenario-Based MPC for Energy-Efficient Building Climate Control under Weather and Occupancy Uncertainty. In *European Control Conference (ECC)*, 1029–1034.

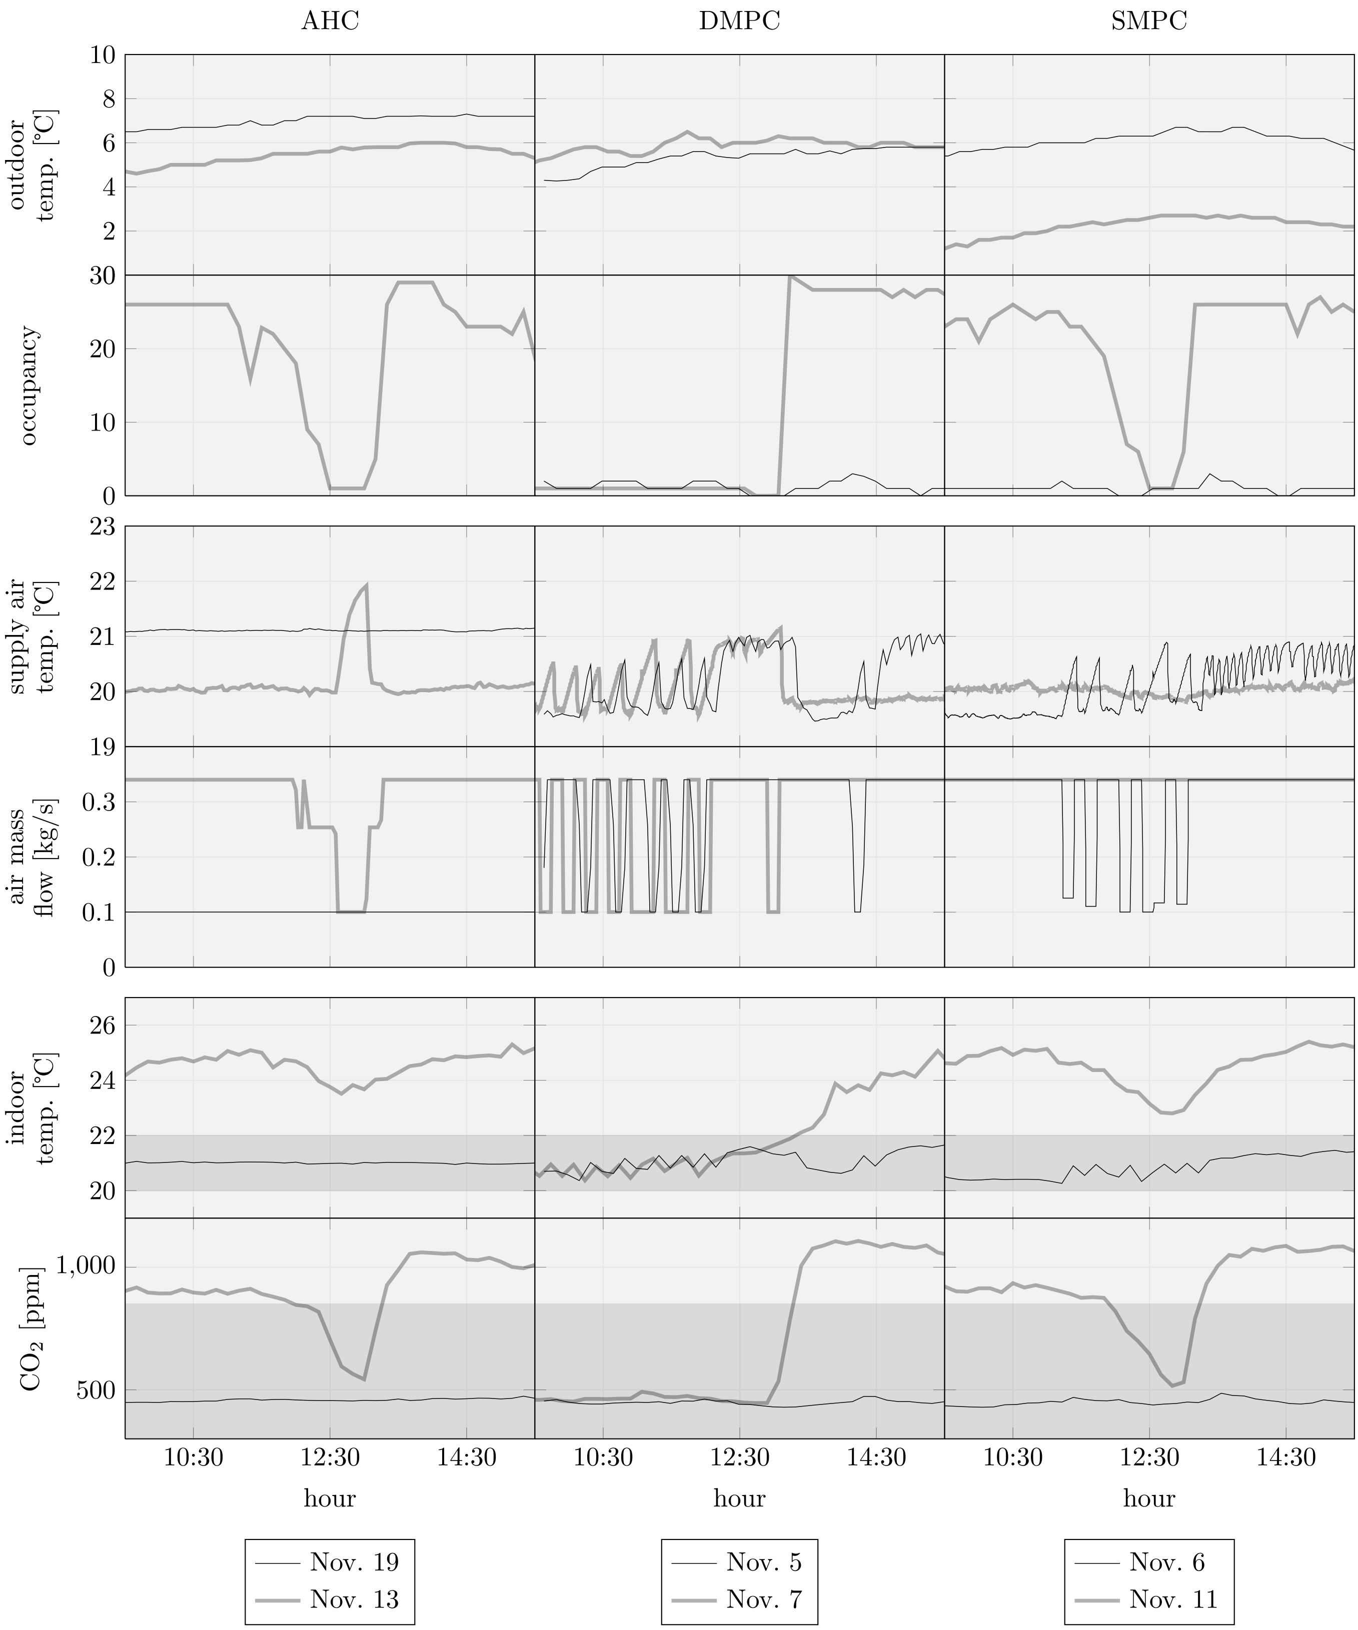

Fig. 6. Disturbances, CO_2 levels, indoor temperatures and control inputs profiles for high- and low-occupancy experimental tests. The shaded areas represent the comfort bounds.