Learning-based predictive control of the cooling system of a large business center

E. Terzi[a], T. Bonetti[a], D. Saccani[a], M. Farina[a], L. Fagiano[a], R. Scattolini[a]

[a] *Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Milan 20133, Italy*
(email: {enrico.terzi,riccardo.scattolini}@polimi.it)

Abstract

This paper describes the design of an advanced control algorithm for the cooling system of a large business and commercial center. This complex system comprises phenomena that are difficult to model with physical principles, such as the demand of the users, heat transport phenomena in a large and complex pipe network, and the behaviour of cooling elements installed by third parties. Motivated by these features, a learning-based model predictive control (MPC) approach is proposed in this paper. The data-driven procedure requires only high-level prior information, making it is easy to implement and to replicate on similar systems. Specifically, to derive a dynamic model of the plant, a comparison among AutoRegeressive eXogenous (ARX), Output Error (OE), Echo State Networks (ESN) and Long Short Term Memory (LSTM) neural networks has been performed. The latter have been eventually selected in view of their higher predictive performance on a validation dataset. Then, an output feedback MPC scheme has been designed to cope with the nonlinearity of the model and the presence of boolean control inputs, corresponding to the *on/off* switching of the cooling units. The resulting MPC algorithm has been tested on a grey-box model of the system, showing significant potential improvements with respect to the baseline controller currently employed.

Keywords: Cooling station, control for energy saving, learning-based predictive control, constrained control, nonlinear control, HVAC systems.

1. Introduction

According to [1], over 38% of the energy consumption in Europe is due to buildings and, within that, about 76% pertains to Heating Ventilation and Air Conditioning (HVAC) systems [2]. This motivates the research efforts aimed at improving the efficiency of HVACs by means of sophisticated modelling, simulation and control strategies. Models and simulation tools of HVAC systems can be based on black-box identification or physical equations, see the reviews [3], [4], [5], [6] and the references therein. As for the design of efficient control systems, several control strategies have been investigated over the years,

starting from classical controllers with Proportional-Integral-Derivative (PID) action [7], to more recent research applying agent-based control systems [8], [9], adaptive fuzzy algorithms [10], [11], [12], or artificial intelligence methods as reviewed in [13]. Model Predictive Control (MPC) has been widely studied in view of the possibility to include into the problem formulation state and input constraints, to consider nonlinear systems, to cope with different goals, like temperature tracking and power consumption minimization, and to take advantage of the available disturbance predictions, such as weather forecasts or building occupancy, see the review papers [14], [15], [16], [17]. However, an important drawback of MPC for HVAC control is the need of a mathematical description of the thermal behaviour of the building: if done with first-principles, such a modeling task is difficult and hardly portable from one building to another [13].
In this paper, we consider the design of an advanced control algorithm for the cooling station of a large business and commercial site located in Milan and composed of five buildings with fifteen floors each, see Figure 1. The buildings contain offices and commercial spaces, two server rooms, a canteen, an auditorium, and other facilities. About 7000 people work in the center Monday to Friday, and about 2000 people are present on Saturday, while on Sunday the main thermal load is due to the servers only. The cooling station, consisting of four chillers and one absorber, provides cold water to the offices air-conditioning in spring and summer, while it disperses the heat produced by the servers and the data centers over the whole year.
A semi-physical (i.e., grey-box) model of this system, based on individual models of the system's devices, i.e. chillers, cooling towers, pipes, absorber, has been developed in [18] and [19], see also [20] for a detailed derivation of the overall model. These models have been obtained as a mix of physical (i.e., mass and energy balance) equations and black-box components identified based on available plant data. Then, the individual models of the devices have been assembled to obtain the overall plant model, subsequently used to re-tune the relay-based control system currently implemented. We refer to such a control system as "baseline" controller.
The grey-box model presented in [18] and [19] well describes the plant dynamics, however, its derivation and tuning turned out to be a very time consuming task. Moreover, in many cases the information required to derive such a model might not be available, due to, for example, changes that occurred to the plant over the years or installation of components from third parties. In these cases, retrieving such information might require even longer time and additional costs. For this reason, in this paper we propose an alternative, fully data-driven approach, which requires only very high-level prior information and is therefore easily applicable to design MPC algorithms for complex HVAC systems. At first, different black box dynamic models with linear and nonlinear structures have been estimated, in order to select the best approach for the problem at hand. Specifically, in the linear case Autoregressive eXogenous (ARX) and Output Error (OE) models have been considered [21], while in the nonlinear case ECHO state [22], [23], and Long Short Term Memory (LSTM) [24] recurrent neural networks have been tested. Among the estimated models, LSTM proved

to be the most effective ones and have been used in the subsequent control design phase.
The MPC controller has been designed based on the estimated nonlinear LSTM model. In order to reduce the computational burden required to solve the underlying optimization problem in each sampling period, successive linearization along the predicted state trajectories has been adopted, inspired by the algorithm in [25]. In this way, and by considering a quadratic cost function, the optimization problem to be solved on line turns out to be an Integer Quadratic Programming (IQP) one, since the control variables are boolean, and represent the *on/off* switching of the chillers. The designed MPC, used to control the semi-physical model, achieves good results in terms of constraints satisfaction and reduced power consumption (-14%) with respect to the baseline controller.

Figure 1: Business/Commercial area considered in this paper

The main contributions we provide in the present paper are:

- The application of a recent and powerful class of recurrent Neural Networks, i.e. LSTM networks, as a dynamical model of a real complex system, starting from the dataset sampled in its routine operations.
- The design of a novel predictive control scheme embedding the derived LSTM model in the controller, and endowed with advanced features such as the Extended Kalman filter and the application of a linearization algorithm to lighten the computational burden.
- The successful test of the designed control scheme embedding the learned model on a semi-physical model of the same system, described in previous contributions and proven to accurately reproduce the real system behavior. In this way we considered also a mismatch between the model (i.e. LSTM)

and the plant (i.e. the semi-physical model), that is unavoidable in real applications.

The whole learning-based algorithm requires only high-level information about the plant, i.e. the involved variables, namely the inputs (active chillers), disturbances (external conditions, waste heat available to the absorber) and outputs (delivery and return temperatures, and absorbed energy). This, together with the relatively easy training of LSTM networks, eases the applicability of this approach to similar plants, where normally these quantities are already logged by standard control systems and/or devices. The paper is organized as follows. In Section 2 the cooling station is described. The adopted identification procedures are described and compared in Section 3. Section 4 presents the control algorithm and its application to the grey-box model is shown in Section 5. Finally, conclusions and hints for future developments are discussed in Section 6.

Notation Given a vector $v \in \mathbb{R}^n$, we denote $v.^2$ the vector obtained by squaring vector v element-wise, and $diag(v)$ the diagonal matrix with v on the diagonal. Given a symmetric matrix $Q \succ 0$, we denote $||v||_Q^2 = v^T Q v$. Matrix $0_{a,b}$ is the null matrix of dimension a, b. We denote the Hadamard (element-wise) product with $\circ$.

2. Plant description and control goals

The synoptic of the cooling station is shown in Figure 2. The plant is composed of the following four different subsystems.

Manifolds Two main recirculation loops are present. The *primary loop* includes the delivery manifold, the by-pass valve, and the return manifold, and it is closed across the chillers and the absorption chiller. The by-pass valve acts as a link between the delivery and return manifolds; its role is to guarantee the hydraulic balancing. The water usually flows through the valve from the delivery to the return manifold, however in some situations it may also flow in the opposite way. In the latter case, the temperature measured in the delivery manifold may be different from the one that is actually experienced by the users. The *secondary loop* pumps water from the primary loop to the users, via a group of four modulating pumps, and channels it back to the cooling machines. This loop splits in three collectors leading to different groups of end-users, both on the forward and on the return line. The secondary loop pipes can be very long, so that physical transport delays, solar radiation, and thermal exchange with the environment are significant.

Chillers Chillers are electric cooling machines that remove heat from water and refrigerate it. Each one includes a cooling tower, condensing water pumps, and chilled water distribution pumps. The distribution pumps can operate only according to an *on/off* behaviour. When switched *on*,

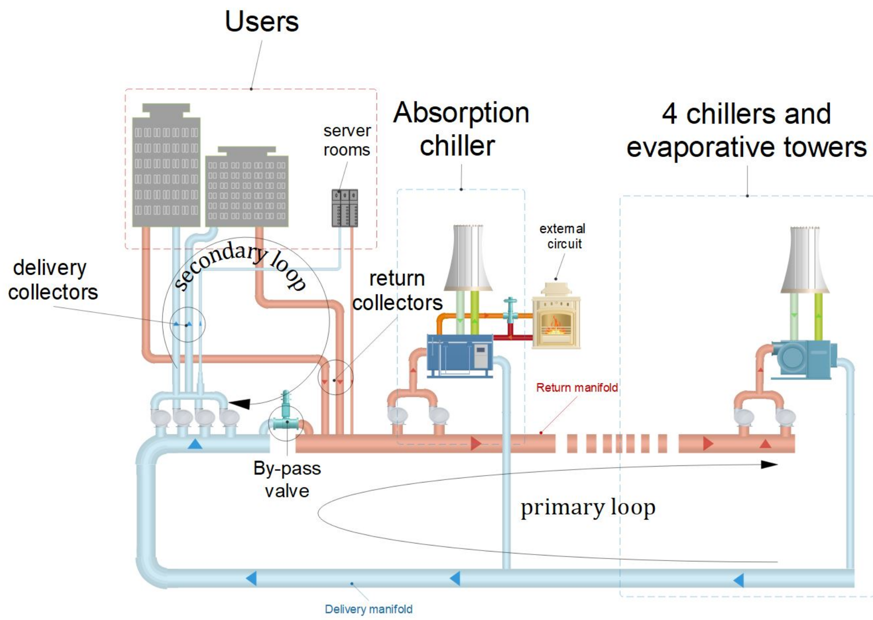

Figure 2: Synoptic scheme of the system under analysis

they circulate water with a constant flow rate, while when they are in the *off* state the water flowing through the chillers is not cooled, and its flow rate depends on the pressure drop between the delivery and the return manifolds. The flow rate in *off* conditions is crucial to guarantee a proper mass balance in the system.

Absorption chiller Absorption chillers, or absorbers, are thermal machines that refrigerate water by exploiting the heat provided by external systems, such as boilers or heat generators, avoiding heat dispersion to the environment, so increasing the overall efficiency of the station. Therefore, absorbers essentially behave as thermal pumps in parallel configuration with the chillers. Their behaviour significantly depends on external variables, and for this reason their modelling is a difficult task.

Users The users consist of offices, meeting rooms, server rooms, shops, etc., in the five buildings and the rooftop. They receive water from the three collectors in the secondary loop and warm it depending on their requests. In total, the users are divided in eight end users, each one having inlet and outlet temperature measures.

2.1. Control goals

The goal of the control system is to regulate the switching pattern of the chillers in order to provide cooled water to the users. This means that the

cooling station must guarantee a balance between the thermal power requested by the users and the one produced by the station itself. The available control variables are the chillers' status, i.e. binary variables s_i, $i = 1, ..., 4$ representing the *on/off* condition of any chiller. As for the absorption chiller, it is considered always *on* when there is waste heat available from the hot circuit; therefore, it is not taken as a decision variable, but it acts as a disturbance in the control system.
The control objectives, sometimes conflicting, are:

- To make the water temperature in the delivery manifold lie in a given range, without exceeding bounds for a long time interval, and stay close to a set-point value of $8.5°C$.
- To avoid the activation of the chillers for short intervals. In fact, the chiller activation for too short time slots entails a cost in terms of wearing of the machine and of power waste, and is not useful to improve the performance of the plant.
- To minimize the number of active chillers (i.e., in *on* state) to reduce the power consumption.
- To switch off the chillers during night and weekends. In fact, in these periods the thermal load requested by the users is small and in many cases the absorption chiller can compensate for it.
- To exploit the absorption chiller as much as possible.
- To balance the use of the four chillers equally activated to help balancing their use.

3. Model identification

3.1. Dataset and model structures

The available experimental data comprise the evolution of hundreds of plant variables along the period May 2016 - October 2016. Data have been collected with sampling time $T_s = 60$ s and in closed-loop working conditions, i.e. under relay-based control (baseline controller) aimed at maintaining the water temperature in the delivery manifold around 8.5^oC. All the data used in the identification procedure have been normalized with respect to their mean and standard deviation.
The identified model shall properly describe the dynamics of the following *output* variables: (i) the temperatures T_{man}, T_{ret} in the delivery and return manifolds, respectively. These models are needed to enforce, in the control design phase, suitable safety and comfort constraints; (ii) the power Q_c absorbed by the chillers, which is essentially the cost to be minimized; (iii) the users' power consumption Q_u. The latter is used to estimate, and possibly predict, the thermal load to be compensated by the chillers.

The available model *inputs* are the number of active chillers, the absorber's power Q_{hot}, the external temperature T_{ext} and humidity H_{ext}, the day of the week, and the time of the day.

According to a well-established procedure, see [21], the identification phase has been performed on a subset of the data (identification set), while the model performances have been cross-validated on a further data set (validation set). Four different model structures have been taken into consideration: 1) linear ARX and 2) OE models, see [21]; 3) ECHO state recurrent neural networks, see [22, 26]; 4) LSTM neural networks, see Section 3.2. The accuracy of the estimated models - regarding the dynamics of T_{man} and T_{ret} - has been assessed by computing the corresponding Normalized Root Mean Square Error (NRMSE) on the validation data. The results are reported in Table 1, which clearly show

$NRMSE$				
	ARX	OE	$ECHO$	$LSTM$
T_{man}	0.0758	0.0571	0.0489	0.0398
T_{ret}	0.0726	0.0480	0.0405	0.0349

Table 1: Comparison between $NRMSE$ (validation data) of the identified linear and nonlinear models

that LSTM networks outperform the other model structures. For this reason, LSTM networks have been selected and used to model the overall plant and to design the predictive control algorithm.

3.2. LSTM models

The general structure of the LSTM models used in this work is the following.

$$\begin{cases} x(k+1) = & \sigma_g(W_f\ u(k) + U_f\ \xi(k) + b_f\ ) \circ x(k) + \\ & \sigma_g(W_i\ u(k) + U_i\ \xi(k) + b_i\ ) \circ \sigma_c(W_c\ u(k) + U_c\ \xi(k) + b_c\ ) \\ \xi(k+1) = & \sigma_g(W_o u(k) + U_o \xi(k) + b_o) \circ \sigma_c(x(k+1)) \\ y(k) = & W_{out}\xi(k) + b_{out} \end{cases} \tag{1}$$

In (1), $u \in \mathbb{R}^m$ is the input vector, $[x^T, \xi^T]^T \in \mathbb{R}^n$ is the network state vector, while $y \in \mathbb{R}^p$ is the output vector. The vector $x \in \mathbb{R}^{n_x}$ is also denoted hidden state, while $\xi \in \mathbb{R}^{n_x}$ is denoted output state, where $n_x = n/2$. $(W_f, W_i, W_o, W_c) \in \mathbb{R}^{n_x \times m}$, $W_{out} \in \mathbb{R}^{p \times n_x}$, $(U_f, U_i, U_o, U_c) \in \mathbb{R}^{n_x \times n_x}$ are weighting matrices, $(b_f, b_i, b_o, b_c) \in \mathbb{R}^{n_x}$, $b_{out} \in \mathbb{R}^p$ are biasing vectors, and $\sigma_c(\circ) = \tanh(\circ) \in (-1, 1)$ and $\sigma_g(\circ) = \frac{1}{1+e^{-\circ}} \in (0, 1)$ are the so-called activation functions.
Two different submodels with this structure have been identified separately, as described below.

3.2.1. Model of the delivery and return manifolds temperatures

These models describe the dynamics of the delivery and return temperatures T_{man} and T_{ret}, respectively, inside the manifolds. The inputs include all the external signals affecting the plant, namely the absorber power Q_{hot}, the users' power request Q_u, the external temperature T_{ext}, the external relative humidity H_{ext}, and the state of the chillers. To represent the latter, we use $\sum_{i=1}^{4} s^i$, i.e, the sum of the binary variables representing the chillers' status. Note also that, while the absorber power Q_{hot} is measured, the users' power request Q_u is, in turn, obtained with the estimated model discussed later on in Section 3.3.
A comparison between the real transients of T_{man} and T_{ret}, collected in about three days, and the outputs of the estimated LSTM models is reported in Figure 3.

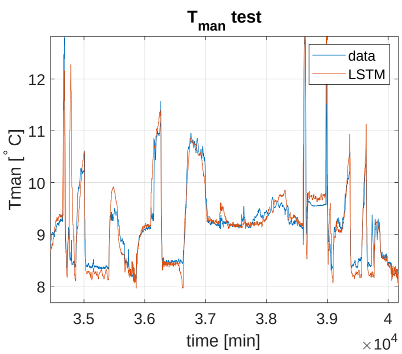

(a)

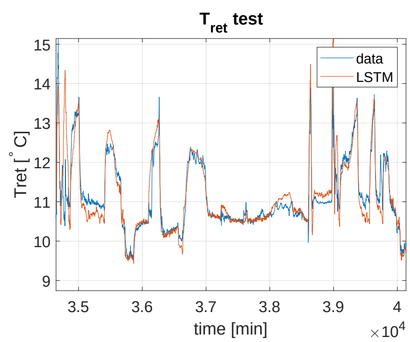

(b)

Figure 3: LSTM with 150 hidden states model - Temperature profile in the delivery manifold (a) and in the return manifold (b).

3.2.2. Chillers' power consumption

An LSTM recurrent network has also been used to model the power Q_c absorbed by the chillers. The inputs of the model are the same already considered for T_{man} and T_{ret}. The NRMSE obtained with the identification and the validation data is reported in Table 2, while Figure 4 shows an example of measured and estimated power in the validation test.

	$NRMSE_{train}$	$NRMSE_{test}$
$\sum_{i=1}^{4} P_{ch}^i$	0.0814	0.0889

Table 2: $NRMSE$ of the LSTM model for absorbed power

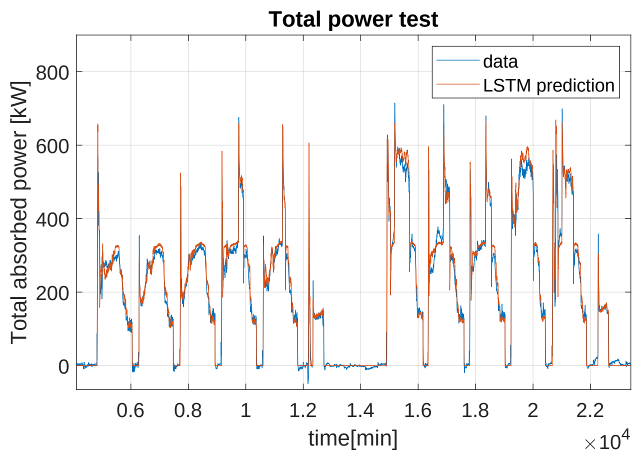

Figure 4: LSTM model, with 200 states, of the power Q_c consumed by the chillers. The figure displays two weeks of data.

3.3. Model of the users' consumption

In this work we use the estimator of the users' consumption Q_u - during the working days of the week (Monday-Friday), and weekends (Saturday and Sunday) - described in [18] and [19]. It consists of two different feedforward networks with one hidden layer. The inputs of these networks are the external humidity and temperature, the day of the week and the time of the day, expressed in minutes.

3.4. Overall model

By interconnecting the above-described black-box models, the overall system model is obtained, presented in Figure 5. Its mathematical equation is of the form:

$$\chi(k+1) = \eta(\chi(k), u(k), d(k)) \quad (2)$$

$$y(k) = g(\chi(k)) \quad (3)$$

The control input $u \in \{0, 1, 2, 3, 4\}$ is the number of active chillers, while $d = [\ Q_{hot} \quad T_{ext} \quad H_{ext} \quad day \quad min\ ]^T$ is the vector of exogenous variables. All the elements of d can be predicted based on the management of the plant (Q_{hot}), the weather forecasts (T_{ext}, H_{ext}), and the considered time instant (day, min). Vector χ is the state of the overall model, while the output

$$y(k) = \begin{bmatrix} y_1(k) \\ y_2(k) \end{bmatrix} = \begin{bmatrix} g_1(\chi(k)) \\ g_2(\chi(k)) \end{bmatrix}$$

is composed of two sub-vectors: measurable outputs $y_1 = [\ T_{man} \quad T_{ret}\ ]^T$ and non-measurable ones $y_2 = [\ Q_c \quad Q_u\ ]^T$.

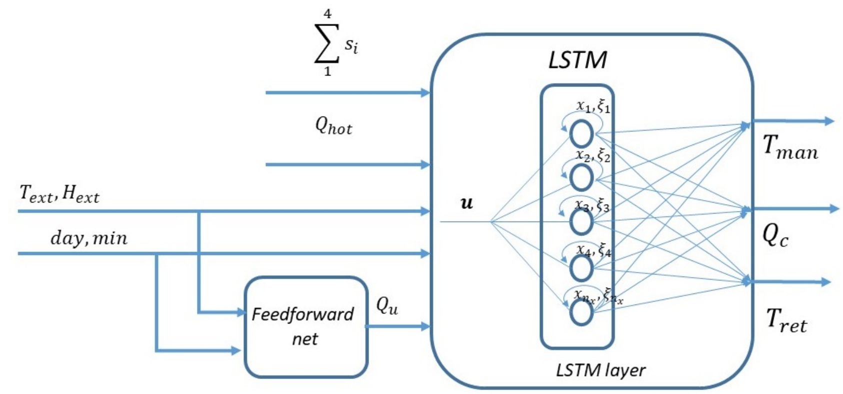

Figure 5: Block scheme of the LSTM model

4. Control design

The control system proposed in this paper is composed of a state estimator and an MPC controller, as described below.

4.1. State estimation

The estimate of the states of system (2) is required to implement the state-feedback MPC algorithm presented in the following. To this end, an Extended Kalman Filter (EKF), able to cope with nonlinear systems, see [27], has been implemented with the following structure:

$$\tilde{\chi}(k+1) = \eta(\tilde{\chi}(k), u(k), d(k)) + L(k)[y_1(k) - \tilde{y}_1(k))] \quad (4)$$

$$\tilde{y}_1(k) = g_1(\tilde{\chi}(k)) \quad (5)$$

where $\tilde{\chi}(k)$ and $\tilde{y}_1(k)$ are the estimates of $\chi(k)$ and $y_1(k)$, respectively. $L(k)$ is the time-varying gain obtained by linearization of the system at the current estimate, see again [27]. Notably, the matrices of the linearized model are relatively straightforward to be computed analytically in view of the fact that (2) is obtained based on the LSTM model structure (1).
Note also that the estimator is fed with a subset of the output vector, i.e. $y_1(k)$, which contains only the measurable variables.

4.2. Successive linearization approach

In MPC, at any time instant k the control law is computed by minimizing a suitable cost function under the constraint given by the dynamics of the system under control, besides state and control limitations. The cost function typically penalizes, over a prediction horizon of length N, the future control variables $u(k+i)$, $i = 0, ..., N-1$, and the future deviations of the output $y(k+i)$, $i = 1, ..., N$, with respect to given reference values. The solution to the

resulting optimization problem may be difficult to compute when the system is nonlinear. An efficient approach to overcome this limitation is to linearize the system's dynamics around the predicted future trajectories, as proposed in [25], thus approximating the system with a Linear Time-Varying (LTV) model. This linearization procedure is described in the following, while the considered optimization problem is formally stated at the end of this section.
At time k, consider a prediction horizon of length N and a candidate input sequence

$$\widehat{U}(k) = \begin{bmatrix} \hat{u}(k)^T & \hat{u}(k+1)^T & \dots & \hat{u}(k+N-2)^T & \hat{u}(k+N-1)^T \end{bmatrix}^T$$

The sequence $\widehat{U}(k)$ can be defined based on the optimal input sequence, $U^\circ(k-1)$, computed at time instant $k-1$

$$U^\circ(k-1) = \begin{bmatrix} u^{\circ T}_{k-1|k-1} & u^{\circ T}_{k|k-1} & \dots & u^{\circ T}_{k+N-2|k-1} \end{bmatrix}^T,$$

where for a generic variable x, $x_{k+j|k}$ denotes its value pertaining to time step $k+j$ predicted at time k. In particular, one can set $\widehat{U}(k)$ as:

$$\widehat{U}(k) = \begin{bmatrix} u^{\circ T}_{k|k-1} & \dots & u^{\circ T}_{k+N-2|k-1} & u^{\circ T}_{k+N-2|k-1} \end{bmatrix}^T$$

where the term $u^\circ_{k+N-2|k-1}$ is considered twice as a reasonable estimate of $u^\circ_{k+N-1|k-1}$. In addition, define the vector of disturbance predictions

$$\widehat{D}(k) = \begin{bmatrix} \hat{d}(k)^T & \hat{d}(k+1)^T & \dots & \hat{d}(k+N-1)^T \end{bmatrix}^T$$

Letting $\hat{\chi}(k)$ be the actual state estimate and considering the trajectory generated using $\widehat{U}(k)$ we can compute the predicted state $\hat{\chi}(k+i)$, $i = 1, \dots, N$ according to

$$\hat{\chi}\,(k+i+1) = \eta\left(\hat{\chi}(k+i), \hat{u}(k+i), \hat{d}(k+i)\right), \qquad i = 0,\ 1, \dots,\ N-1 \quad (6)$$

At this point we can define the linearized system matrices as

$$\begin{aligned} A_{k+i} &= \left.\frac{\partial \eta}{\partial \chi}\right|_{\hat{\chi}(k+i), \hat{u}(k+i),\ \hat{d}(k+i)} \\ B_{k+i} &= \left.\frac{\partial \eta}{\partial u}\right|_{\hat{\chi}(k+i), \hat{u}(k+i),\ \hat{d}(k+i)} \\ M_{k+i} &= \left.\frac{\partial \eta}{\partial d}\right|_{\hat{\chi}(k+i), \hat{u}(k+i),\ \hat{d}(k+i)} \end{aligned} \quad (7)$$

resulting in the linear time-varying system (see [20]):

$$\chi(k+i+1) = A_{k+i}\chi(k+i) + B_{k+i}u\,(k+i) + M_{k+i}\hat{d}(k+i) + h(k+i) \quad (8)$$

where the additive term

$$h\,(k+i) = \hat{\chi}(k+i+1) - (A_{k+i}\hat{\chi}(k+i) + B_{k+i}\hat{u}(k+i) + M_{k+i}\hat{d}(k+i)) \quad (9)$$

can be computed at time k over the whole horizon $i = 1, \dots, N-1$.

4.3. Reference input trajectory definition

In this section we describe how the reference input trajectory, to be used in the cost function, is defined. The computation is based on the following available trajectories: (*i*) the prediction of the absorber's available power $Q_{hot}(k+i)$, $i = 0, \ldots, N-1$; (*ii*) the prediction of the users' power request $Q_u(k+i)$, $i = 0, \ldots, N-1$; (*iii*) the nominal power absorbed by each chiller Q_{cu}. Exploiting this information, we predict the number of chillers required to be active during the prediction horizon as

$$u_r(k+i) = \text{round}\left(\frac{Q_u(k+i) - Q_{hot}(k+i)}{Q_{cu}}\right) \tag{10}$$

In short, we define

$$U_r(k) = \begin{bmatrix} u_r^T(k) & u_r^T(k+1) & \ldots & u_r^T(k+N-1) \end{bmatrix}^T \tag{11}$$

4.4. Disturbance compensation

The state estimate provided by the described EKF at each time instant is exploited at every optimization occurrence as initial state to compute the predicted trajectories thanks to the estimated LSTM model. This estimate suffers naturally of model/plant mismatch given the intrinsic difference between the real plant (simulator) and the LSTM model. In order to partially compensate it, according to [28], we provide the controller with a signal $e(k) = y_1(k) - \tilde{y}_1(k)$ that is added as a biasing term to the future output predictions.

4.5. Optimization problem and MPC algorithm

The MPC problem has been defined by choosing the prediction horizon $N = 120$, corresponding to two hours of time, and long enough to include the dominant dynamics of the system. Moreover, in order to reduce the overall computational burden associated with the solution of the underlying optimization problem, it has been assumed that the control variable can vary only every $M = 20$ minutes, so that the total number of variations in the considered horizon is $N_s = N/M = 6$. This choice is motivated also by the cooling dynamics of the chillers, which settles in about 10 min, and can not be switched so frequently.

The vector of optimization variables is defined as

$$\bar{U}(k) = \begin{bmatrix} u(k)^T & u(k+M)^T & \ldots & u(k+(N_s-1)M)^T \end{bmatrix}^T \tag{12}$$

and the control variables are defined at any time instant as

$$u(k+i) = u(k + \left\lfloor \frac{i}{M} \right\rfloor M)\,,\ i = 0, ..., N-1 \tag{13}$$

and collected in vector

$$U(k) = \begin{bmatrix} u(k)^T & u(k+1)^T & \ldots & u(k+N-1)^T \end{bmatrix}^T \tag{14}$$

In addition, as customary in MPC, the vector of control increments

$$\delta\bar{U}(k) = \bar{U}(k) - \bar{U}(k-M) \tag{15}$$

is weighted in the cost function to reduce the variability of the MPC control action over time, and consequently limit the switches of the chillers. Moreover, letting

$$\tilde{U}(k) = \left[\begin{array}{cccc} u^{o^T}_{k|k-M} & u^{o^T}_{k+M|k-M} & \cdots & u^{o^T}_{k+(N_s-2)M|k-M} \end{array} \right]^T \tag{16}$$

it is also worth penalizing the vector difference

$$\delta\tilde{U}(k) = \bar{U}(k) - \tilde{U}(k) \tag{17}$$

to reduce the effects of the linearization error due to the procedure described in Section 4.2. We define the difference

$$\delta U_r(k) = U(k) - U_r(k) \tag{18}$$

which can be included in the cost function to force the future control action to be close to the predicted ones based on the future expected evolution of the external environmental and load conditions.

Denoting with $\Lambda(k)$ a vector of nonnegative slack variables introduced to guarantee the feasibility of the problem at any time instant, see [29], the optimization problem is:

$$\min_{\bar{U}(k),\Lambda(k)} J = ||\bar{U}(k)||^2_R + ||\delta\bar{U}(k)||^2_{\Delta R} + ||\delta\tilde{U}(k)||^2_{\Delta\tilde{R}} + ||\delta U_r(k)||^2_{\Delta R_r} + \rho||\Lambda(k)||_1 \tag{19}$$

The minimization of (19) at any time instant k must be performed subject to the following set of constraints, where inequalities between vectors are element-wise:

$$u(k+i) = u(k + \left\lfloor \frac{i}{M} \right\rfloor M)\,,\; i = 0, ..., N-1 \tag{20a}$$

$$\chi^*(k) = \tilde{\chi}(k) \tag{20b}$$

$$\hat{\chi}(k) = \tilde{\chi}(k) \tag{20c}$$

$$\chi^*(k+i+1) = A_{k+i}\chi^*(k+i) + B_{k+i}u\,(k+i) + M_{k+i}\hat{d}(k+i) + h(k+i),\quad i = 0, ..., N-1 \tag{20d}$$

$$\hat{\chi}\,(k+i+1) = \eta\left(\hat{\chi}(k+i), \hat{u}(k+i), \hat{d}(k+i)\right),\; i = 0,\; 1, \ldots,\; N-1 \tag{20e}$$

$$h\,(k+i) = \hat{\chi}(k+i+1) - (A_{k+i}\hat{\chi}(k+i) + B_{k+i}\hat{u}(k+i) + M_{k+i}\hat{d}(k+i)),\quad i = 0, ..., N-1 \tag{20f}$$

$$\tilde{\chi}(k+1) = \eta(\tilde{\chi}(k), u(k), d(k)) + L(k)[y(k) - g(\tilde{\chi}(k))] \tag{20g}$$

$$y_1^*(k+i) = g_1(\chi^*(k+i)) + e(k),\; i = 0, ..., N-1 \tag{20h}$$

$$Y_{min} - \Lambda(k) \le y_1^*(k+i) \le Y_{max} + \Lambda(k),\; i = 0, ..., N-1 \tag{20i}$$

$$0 \le U(k) \le 4 \tag{20j}$$

$$-1 \le \delta\bar{U}(k) \le 1 \tag{20k}$$

$$-1 \le \delta\tilde{U}(k) \le 1 \tag{20l}$$

$$\Lambda(k) \ge 0 \tag{20m}$$

Where the matrices A_{k+i}, B_{k+i}, M_{k+i} are defined in (7), the positive definite diagonal matrices R, ΔR, $\Delta\tilde{R} \in \mathbb{R}^{N_s m \times N_s m}$ and $\Delta R_r \in \mathbb{R}^{Nm \times Nm}$ in (19) are the tunable parameters of the controller, while ρ is a sufficiently large weight so that $\Lambda(k)$ is close to zero when a feasible solution exists.
The effects of these parameters on the control tuning can be summarized as follows:

- Larger values of R elements induces the controller to decrease the overall number of active chillers;
- Larger values of ΔR elements induces the controller to reduce the overall number of switches;
- Larger values of $\Delta\tilde{R}$ elements induces the controller to reduce the deviation of the optimal solution from the candidate solution $\hat{U}(k)$ used to generate the state trajectory for linearization;
- Larger values of ΔR_r elements induces the controller to reduce the deviation of the optimal solution from the reference trajectory $U_r(k)$;
- ρ is a penalty parameter to force the slack variables $\Lambda(k)$ to be close to zero when a feasible solution exists.

In (20i), $Y_{min} = \begin{bmatrix} T_{manMin} & T_{retMin} \end{bmatrix}^T$ and $Y_{max} = \begin{bmatrix} T_{manMax} & T_{retMax} \end{bmatrix}^T$ specify the minimum and maximum values allowable to the temperatures inside

the manifolds. The constraints on $\delta\bar{U}(k)$ are included to force the control variation at successive time instants to be at most equal to one, while the constraints on $\delta\tilde{U}(k)$ are introduced to force the real control variable to be close to the one used in the linearization procedure.
Once the optimal control vector $\bar{U}(k)$ has been computed, and according to a standard receding horizon principle, only its first element $u^o_{k|k}$ is applied and the overall optimization procedure is repeated M time instants later.

Overall, the complete control scheme is presented in Figure 6, where the process is represented by the physical simulator described in [20].

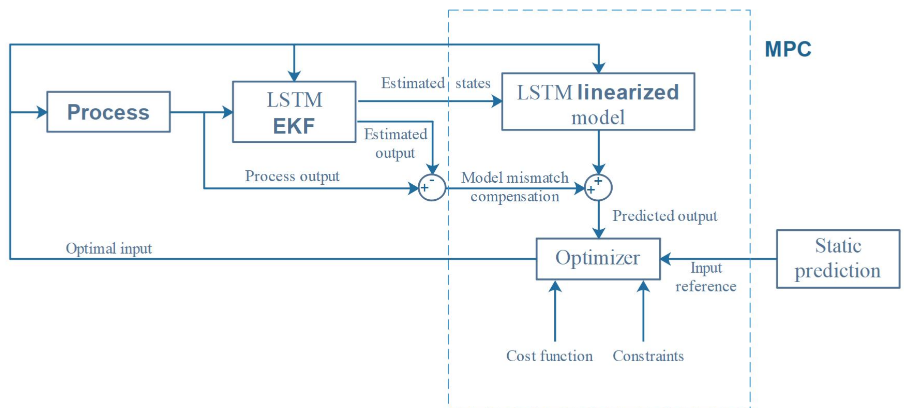

Figure 6: Control scheme of the plant

5. Numerical results

The presented control scheme has been tested in simulation. The design parameters are reported in Table 3.

$Parameters$	$Value$
ρ	10^9
R	$diag(15)$
ΔR	$diag(55)$
$\Delta\tilde{R}$	$diag(60)$
ΔR_r	$diag(3)$
Y_{max}	$[12 \quad 15]$
Y_{min}	$[5 \quad 5]$

Table 3: Weights of the MPC on equations with static prediction

The signal $U_r(k)$ provided by the static prediction is shown in Figure 7a. Note that the chillers are *off* while the requested power is below a threshold, corresponding to the value that the absorber is able to satisfy, thus fulfilling the fourth goal introduced in Section 2.1. Nevertheless, increasing the magnitude of the weighting matrix ΔR_r to keep the control action close to $U_r(k)$ may still lead to a non-desired switching behaviour. Therefore, a further condition is imposed on the value of the thermal load requested by the users (shown in Figure 7b), in order to decide whether to switch the chillers *off*. Specifically, if $u_r(k) = 0$ and the average value of the forecast thermal load over the prediction horizon is lower than a predefined threshold, then the optimization problem is not solved and the chillers are shut down until one of the two conditions is no longer met. This simple logic is particularly helpful at night, when the thermal request is low, in order to automatically deactivate the controller, other than doing it manually or with a time-based logic.

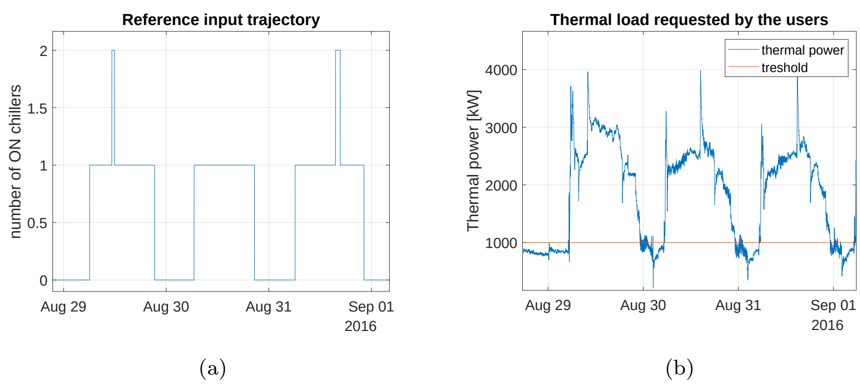

Figure 7: Reference value of the control action provided by the static prediction (a), thermal load requested by the users (b)

To validate the controller performance, a comparison among the proposed solution, the baseline controller actually implemented on the plant, and its optimized version described in [19] has been made. The optimized control logic has been derived exploiting a grey-box system model. The parameters of the existing controller are optimized via nonlinear programming, by minimizing the predicted energy consumption while satisfying the cooling demand. A comparison in terms of water temperature in the delivery manifold, control action and electrical power consumption is reported in Figures 8, 9, 10. It is worth noticing how the baseline controllers, both in their current and optimized versions, are not able to anticipate the user's power request and to satisfy the constraints on the water temperature. These lacks lead to the turn on of the second chiller in the first hours of the day and may cause possible discomforts to the users.
On the contrary, the predictive controller avoids an undesired behaviour of the control variable (e.g., unnecessary switching), exploiting the disturbances forecast, and fulfills the output constraints. The whole multi-objective problem can

be managed by a suitable choice of just few parameters, shown in Table 3.

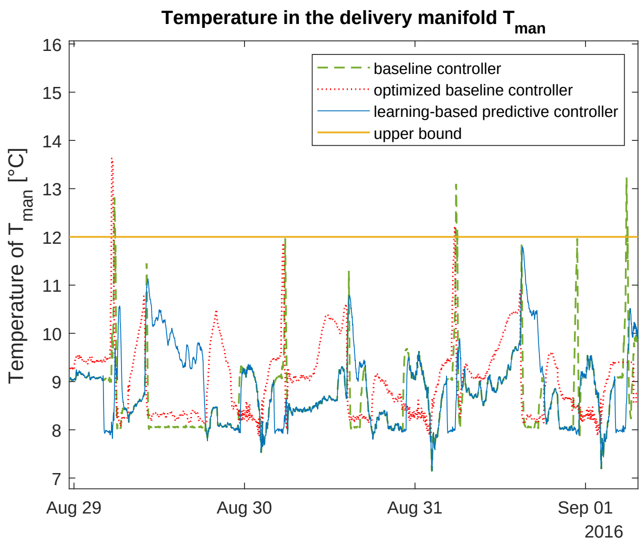

Figure 8: Water temperature in the delivery manifold: comparison among baseline (dashed green line), optimized baseline (dotted orange line) and learning-based predictive controller (solid blue line)

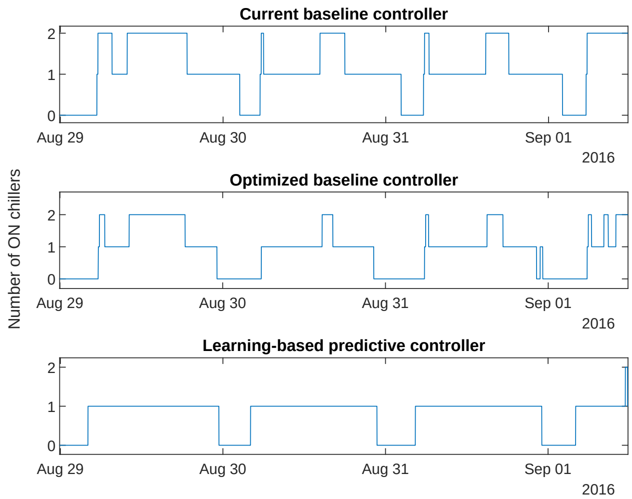

Figure 9: Control action: comparison among baseline, optimized baseline and learning-based predictive controller

It is noticeable that the control scheme achieves a good trade-off between the control objectives described in Section II, in particular:

- The water temperature in the delivery manifold satisfies the upper bound of 12^oC and its average value is similar to the one of measured data with the baseline controller, even though the total number of chillers used is reduced. Specifically, the average value of the simulation data with MPC controller is $0.14°C$ lower than the one resulting from measured data.
- The switching behaviour is reduced; particular attention is given to the early morning, where the prediction of the thermal load allows the control system to anticipate the activation of a chiller and avoids the use of more than one chiller for a short time interval.
- According to the previous comments, the shut-down of the chillers during night is anticipated based on the thermal load prediction.
- The absorption chiller is fully exploited since its cooling capability is expressed in the static prediction formula (see 10) allowing, when possible, to switch the chillers *off*.
- The total power absorbed by the cooling station is reduced, also due to the limitation of the number of active chillers. Using the power model based on LSTM network described in Section 3, a comparison among the power consumption is reported in Figure 10. Specifically, in the period considered, the power consumption, estimated by means of the tailored LSTM model, is decreased from 21250 kWh with the baseline controller to 18260 kWh with MPC controller, thus obtaining a reduction of 14%.

6. Conclusions

This paper has described a procedure for the design of a control system for a large cooling station of a commercial center. A dynamic model of the plant has been obtained with Long Short Term Memory networks trained with experimental data, an Extended Kalman Filter has been used as state estimator, and a Model Predictive Control algorithm has been designed to compute the control variables, which correspond to the *on/off* commands to the chillers. The use of black-box (LSTM) models is a fundamental step in the design procedure, since it allows one to avoid the time consuming physical modeling of the system, made difficult by many uncertain aspects, such as the length of the pipes, transport delays, and chillers' functioning in *off* conditions. The proposed approach is applicable to a wide class of HVAC problems independently from the plant topology and from the nature (boolean or continuos) of the control inputs. Indeed, the use of only high-level information makes this algorithm attractive to every dynamical systems where the physical modelling is too difficult, expensive or time consuming.
The performance provided by the control system, tested in simulation, is very

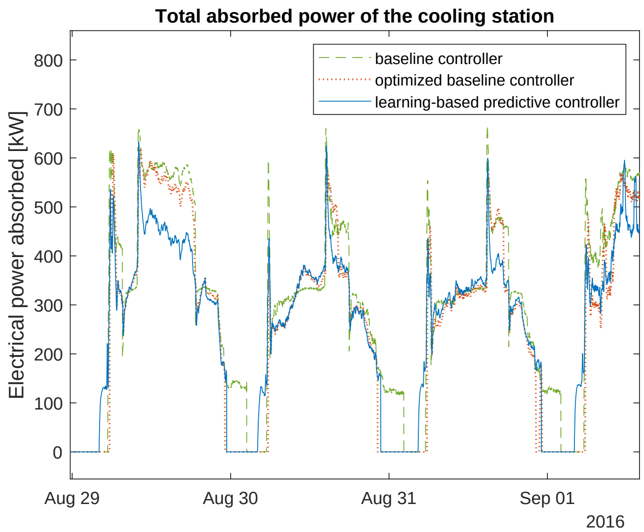

Figure 10: Power consumption: comparison between baseline (dashed green line), optimized baseline (dotted orange line) and learning-based predictive controller (solid blue line)

promising, and shows that the main temperatures of the cooling system are maintained in prescribed limits and significant energy savings are achieved with respect to the relay-based controller currently used.

References

[1] A. Allouhi, Y. El Fouih, T. Kousksou, A. Jamil, Y. Zeraouli, Y. Mourad, Energy consumption and efficiency in buildings: current status and future trends, Journal of Cleaner Production 109 (2015) 118–130.

[2] F. W. Payne, J. J. McGowan, Energy management and control systems handbook, Fairmont Press, ISBN 9780881730289, 1988.

[3] A. Afram, F. Janabi-Sharifi, Review of modeling methods for HVAC systems, Applied Thermal Engineering 67 (1-2) (2014) 507–519.

[4] R. Homod, Review on the HVAC System Modeling Types and the Shortcomings of Their Application, Journal of Energy 2013, doi: 10.1155/2013/768632.

[5] M. Trčka, J. L. M. Hensen, Overview of HVAC system simulation, Automation in Construction 19 (2) (2010) 93–99.

[6] Z. Afroz, G. M. Shafiullah, T. Urmee, G. Higgins, Modeling techniques used in building HVAC control systems: A review, Renewable and Sustainable Energy Reviews 83 (2018) 64–84.

[7] G. Levermore, Building energy management systems: An application to heating, natural ventilation, lighting and occupant satisfaction, Routledge, 2013.

[8] I. T. Michailidis, T. Schild, R. Sangi, P. Michailidis, C. Korkas, J. Fütterer, D. Müller, E. B. Kosmatopoulos, Energy-efficient HVAC management using cooperative, self-trained, control agents: A real-life German building case study, Applied energy 211 (2018) 113–125.

[9] W. Jung, F. Jazizadeh, Comparative assessment of HVAC control strategies using personal thermal comfort and sensitivity models, Building and Environment 158 (2019) 104–119.

[10] S. Hu, C. Hoare, P. Raftery, J. O'Donnell, Environmental and energy performance assessment of buildings using scenario modelling and fuzzy analytic network process, Applied Energy 255 (2019) 113788.

[11] K. Sun, S. Mou, J. Qiu, T. Wang, H. Gao, Adaptive Fuzzy Control for Nontriangular Structural Stochastic Switched Nonlinear Systems With Full State Constraints, IEEE Transactions on Fuzzy Systems 27 (8) (2019) 1587–1601.

[12] J. Qiu, K. Sun, T. Wang, H. Gao, Observer-Based Fuzzy Adaptive Event-Triggered Control for Pure-Feedback Nonlinear Systems With Prescribed Performance, IEEE Transactions on Fuzzy Systems 27 (11) (2019) 2152–2162.

[13] A. I. Dounis, C. Caraiscos, Advanced control systems engineering for energy and comfort management in a building environment—A review, Renewable and Sustainable Energy Reviews 13 (6-7) (2009) 1246–1261.

[14] A. Afram, F. Janabi-Sharifi, Theory and applications of HVAC control systems–A review of model predictive control (MPC), Building and Environment 72 (2014) 343–355.

[15] A. Afram, F. Janabi-Sharifi, A. S. Fung, K. Raahemifar, Artificial neural network (ANN) based model predictive control (MPC) and optimization of HVAC systems: A state of the art review and case study of a residential HVAC system, Energy and Buildings 141 (2017) 96–113.

[16] S. Wang, Z. Ma, Supervisory and optimal control of building HVAC systems: A review, HVAC&R Research 14 (1) (2008) 3–32.

[17] R. Kwadzogah, M. Zhou, S. Li, Model predictive control for HVAC systems—A review, in: 2013 IEEE International Conference on Automation Science and Engineering (CASE), IEEE, 442–447, 2013.

[18] E. Terzi, L. Fagiano, M. Farina, R. Scattolini, Identification of the cooling system of a large business center, IFAC-PapersOnLine 51 (15) (2018) 174–179.

[19] E. Terzi, L. Fagiano, M. Farina, R. Scattolini, Structured modelling from data and optimal control of the cooling system of a large business center, Journal of Building Engineering 28 (2020) 101043.

[20] T. Bonetti, D. Saccani, Modeling and learning-based hybrid predictive control with recurrent neural networks of the cooling station of a large business center, Tech. Rep., MS thesis Politecnico di Milano, 2019.

[21] L. Ljung, System identification, in: Signal analysis and prediction, Springer, 163–173, 1998.

[22] H. Jaeger, Adaptive nonlinear system identification with echo state networks, in: Advances in neural information processing systems, 609–616, 2003.

[23] H. Jaeger, The "echo state" approach to analysing and training recurrent neural networks-with an erratum note, Bonn, Germany: German National Research Center for Information Technology GMD Technical Report 148 (34) (2001) 13.

[24] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural computation 9 (8) (1997) 1735–1780.

[25] P. Falcone, F. Borrelli, H. E. Tseng, J. Asgari, D. Hrovat, Linear time-varying model predictive control and its application to active steering systems: Stability analysis and experimental validation, International Journal of Robust and Nonlinear Control: IFAC-Affiliated Journal 18 (8) (2008) 862–875.

[26] L. B. Armenio, E. Terzi, M. Farina, R. Scattolini, Model Predictive Control Design for Dynamical Systems Learned by Echo State Networks, IEEE Control Systems Letters 3 (4) (2019) 1044–1049.

[27] B. D. O. Anderson, J. B. Moore, Optimal filtering, Courier Corporation, 2012.

[28] M. Morari, U. Maeder, Nonlinear offset-free model predictive control, Automatica 48 (9) (2012) 2059–2067.

[29] J. M. Maciejowski, Predictive control with constraints, Pearson education, 2002.