MITSUBISHI ELECTRIC RESEARCH LABORATORIES
https://www.merl.com

A Comparison of MPC Architectures for a Heat Pump

Bortoff, Scott A.; Quah, Titus; Rawlings, James B.

TR2026-060 May 28, 2026

Abstract

Modern heat pumps are strongly interactive multivariable systems, requiring rigorous application of multi- variable control theory. Important factors to consider are robustness with respect to plant uncertainty, and enforcement of process constraints. This paper summarizes the design of two Model Predictive Control (MPC) architectures for a heat pump application, and compares their design methodologies and performance to a conventional selector-based multi-loop PID architecture. One MPC architecture is an H-infinity loop- shaped MPC that uses inverse optimal control to realize the ro- bustifying compensator as a constrained optimization to enforce constraints. The second is an offset-free MPC architecture that preserves output tracking despite plant-model mismatch and unmeasured disturbances and can be tuned using operational data. We compare the two MPCs and the PID on a constrained heat pump model, assessing closed-loop transients, robustness margins, and tuning complexity. In simulation, all controllers track setpoints with similar performance. During a sensible- heat disturbance, a temperature limit prevents full rejection; even so, both MPC designs counterintuitively reduce the out- door fan speed – which, through coupled system interactions, increases net heat removal and yields improved disturbance rejection without violating constraints. As for robustness, all three designs meet disk-margin targets. For tuning complexity, offset-free MPC > H-infinity loop-shaped MPC > PID. Overall, on this plant, robustness is a tie; offset-free MPC delivers the strongest constrained disturbance handling but requires the most tuning, H-infinity loop-shaped MPC is the middle ground, and PID is simplest.

American Control Conference (ACC) 2026

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

A Comparison of MPC Architectures for a Heat Pump

Titus Quah,[1] Scott A. Bortoff [2] and James B. Rawlings[1]

Abstract— **Modern heat pumps are strongly interactive multivariable systems, requiring rigorous application of multivariable control theory. Important factors to consider are robustness with respect to plant uncertainty, and enforcement of process constraints. This paper summarizes the design of two Model Predictive Control (MPC) architectures for a heat pump application, and compares their design methodologies and performance to a conventional selector-based multi-loop PID architecture. One MPC architecture is an H-infinity loop-shaped MPC that uses inverse optimal control to realize the robustifying compensator as a constrained optimization to enforce constraints. The second is an offset-free MPC architecture that preserves output tracking despite plant-model mismatch and unmeasured disturbances and can be tuned using operational data. We compare the two MPCs and the PID on a constrained heat pump model, assessing closed-loop transients, robustness margins, and tuning complexity. In simulation, all controllers track setpoints with similar performance. During a sensible-heat disturbance, a temperature limit prevents full rejection; even so, both MPC designs counterintuitively reduce the outdoor fan speed – which, through coupled system interactions, increases net heat removal and yields improved disturbance rejection without violating constraints. As for robustness, all three designs meet disk-margin targets. For tuning complexity, offset-free MPC > H-infinity loop-shaped MPC > PID. Overall, on this plant, robustness is a tie; offset-free MPC delivers the strongest constrained disturbance handling but requires the most tuning, H-infinity loop-shaped MPC is the middle ground, and PID is simplest.**

I. Introduction

Modern heat pumps are based on a vapor compression cycle that is actuated by a variable speed compressor, one or more variable speed fans, and one or more electronically actuated expansion valves. Their dynamics are multivariable and interactive, demanding rigorous application of multivariable control theory. Arguably the two most important factors to consider in control design are robustness with respect to plant uncertainty, and enforcement of process constraints.

Driven by market demands, heat pumps must operate over an increasingly large operating envelope over which the dynamics change [1]. Furthermore, heat pumps must operate reliably in a diversity of applications. Hence the plant model representing an installed heat pump varies from one application to another, due to varying pipe lengths and geometries, varying refrigerant charge, and dynamic coupling with the dynamics associated with each particular application. Heat pump controllers must also enforce constraints on actuator limits and a number of process variables (output constraints), often for long periods of time (even steady-state) in order to maintain proper operation. These constraints represent requirements on the refrigerant state at various points in the vapor compression cycle, in order to prevent liquid refrigerant from being ingested into the compressor, or to prevent frost build up on a coil, for example.

Decentralized multivariable feedback control with PID compensation, sometimes gain scheduled, along with selector logic to enforce process constraints is commonly used in production heat pumps. Yet these designs can be difficult to transfer across applications, often requiring extensive retuning and gain scheduling as dynamics change with operating condition, refrigerant charge, and equipment configuration. Moreover, as the numbers of inputs, outputs, and constraints grow, the associated selector and override logic becomes increasingly complex to design, validate, and maintain. Given the wide variety of indoor and outdoor unit combinations offered by manufacturers, a more unified control development process and architecture is desirable.

This motivates consideration of Model Predictive Control (MPC) as a unified framework for multivariable control and constraint enforcement across applications. While prior work has demonstrated the applicability of MPC to vapor-compression and heat pump systems [2], [3], [4], less attention has been given to a head-to-head comparison between MPC designs and the selector-based PID architectures that remain standard in practice, especially from the standpoint of robustness and tuning complexity. Accordingly, this paper compares $\mathcal{H}_\infty$ loop-shaped MPC and offset-free MPC with a baseline selector-based PID design for a single-zone residential heat pump, evaluating robustness margins, tuning requirements, and transient closed-loop behavior. First, we design an $\mathcal{H}_\infty$ loop-shaped MPC [5], which maximizes the robustness with respect to normalized coprime factor uncertainty in the plant, and provides a methodology to tune compensator parameters based on loop shaping, assuming a model is available. Second, we design an offset-free MPC that eliminates steady-state offset and preserves output tracking despite plant-model mismatch and unmeasured disturbances with nonzero mean, by augmenting the predictor with a disturbance/bias model and estimating that bias online [6], [7]. Beyond its classical offset-free property, recent results show that offset-free MPC is robustly stable to sufficiently small estimate errors, setpoint and disturbance changes, and plant-model mismatch [8]. Moreover, as data collected from heat pumps in operation becomes increasingly available, this framework offers the prospect of tuning the controller state estimator directly from data [9], thereby reducing the tuning burden, although that aspect is beyond the scope of the present work. For each methodology, we compare the two proposed designs with a baseline selector-based PID

[1] University of California, Santa Barbara, CA, USA {titusquah, jbraw}@ucsb.edu, [2]Mitsubishi Electric Research Laboratories, Cambridge, MA, USA bortoff@merl.com

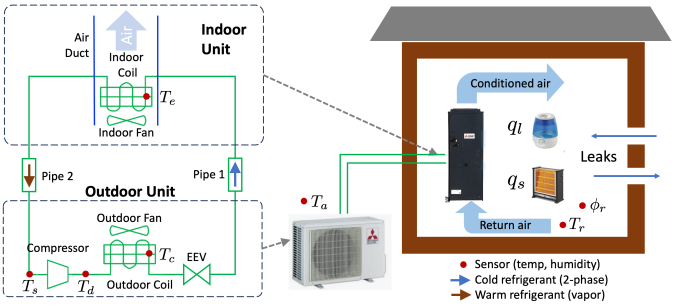

Fig. 1. Heat pump system diagram.

architecture in terms of the number of tuning parameters, methods for their tuning, the resulting robustness margins, and transient performance by simulation.

This paper is organized as follows. In Section II the particular heat pump of interest is described, its model presented, and a set of control requirements is listed. Section III describes the selector-based PID baseline control architecture. Sections IV and V present the $\mathcal{H}_\infty$ loop-shaped MPC, and offset-free MPC designs, respectively, and a simulation comparison, disk margin comparison, and numbers of tunable parameters are compared in Section VI. Some conclusions are drawn in Section VII.

II. System Description & Model

The residential heat pump diagrammed in Fig. 1 consists of an air-cooled outdoor unit (including a variable speed compressor, an outdoor heat exchange coil with variable speed fan, and an electronic expansion valve (EEV)), connected by refrigerant pipes to a ducted indoor air handler unit (including an indoor heat exchange coil with a variable speed fan). The system operates in cooling mode, moving heat from the indoor air to the outdoor air using a vapor compression cycle. A dynamic model of the heat pump coupled to a representative residential building was constructed in the Modelica modeling language [10], [11], [12] using proprietary component libraries [13], [14] and the open source Modelica Buildings Library [15], [16]. The model uses a finite volume method to represent the heat exchangers, and expresses mass, energy and momentum balances associated with the fluid mechanics, heat transfer and phase transitions of refrigerant and moist air. A more detailed description is beyond our scope and we refer the reader to the references.

The model was linearized at a typical operating condition and reduced from a state dimension of 250 to 18 using balanced truncation, modal decomposition and singular perturbation methods. The plant has four control inputs: the compressor speed (CF), the electronic expansion valve counts (EEV), and the indoor and outdoor fan speeds (IFS and OFS). In this work the indoor fan speed is assumed constant. (It may be operated manually, or be under the authority of another controller). The plant model has three disturbance inputs: the ambient (outdoor) temperature T_a, a sensible heat load q_s and a latent heat load q_l. It has seven measured outputs: the room temperature and relative humidity T_r, and ϕ_r, respectively, the compressor discharge temperature T_d, the condensing temperature T_c, the evaporating temperature T_e, the compressor suction temperature T_s, and the outdoor air temperature T_a, summarized in Table I and Fig. 1.

TABLE I

Control Inputs, Disturbances & Measured Outputs.

Name	Description
OFS	Outdoor fan speed (RPM)
IFS	Indoor fan speed (RPM)
CF	Compressor frequency (Hz)
EEV	Electronic expansion valve (counts)
q_s	Sensible indoor heat load disturbance (W)
q_l	Latent indoor heat load disturbance (W)
T_a	Measured outdoor air temperature (°C)
T_r	Measured indoor air temperature (°C)
ϕ_r	Measured indoor air relative humidity (%)
T_d	Measured compressor discharge temperature (°C)
T_c	Measured condensing temperature (°C)
T_e	Measured evaporator temperature (°C)
T_s	Measured suction port temperature (°C)

Fig. 2 shows open-loop step responses from CF (step down 1Hz), LEV (step up 2 counts) and OFS (step up 100RPM) to T_r, T_c, T_d and T_e, showing a two-time scale behavior, with fast dynamics occurring on a 2–3 minute time scale, coupled to slow dynamics occurring on a 24-hour time scale. The slow dynamics are associated with heat transfer from building materials and vary by application, making the low frequency (DC) gain, for example, unknowable.

The requirements for closed-loop control are to regulate the room temperature T_r to a user-specified setpoint r_1, to regulate the evaporator superheat $T_{SH} := T_s - T_e$ to a constant value, typically 5°C, and to regulate the condensing temperature T_c to a reference r_3, typically 40°C. The superheat requirement ensures the refrigerant is a vapor at the compressor inlet, and the condensing temperature requirement is needed at low T_a. Both ensure proper refrigerant distribution around the cycle. Further requirements are: a) Constant values of disturbances (T_a, q_s, q_l) are to be rejected with zero steady-state error; b) Constant references are to be tracked with zero offset; c) Hard constraints on inputs (CF, OFS, EEV) are to be enforced; d) Output constraints $T_e > \bar{T}_e$, $T_d < \bar{T}_d$ and $T_{DSH} > \bar{T}_{DSH}$, where $T_{DSH} = T_d - T_c$, are to be enforced, possibly softly.

III. PI Control Architecture

Fig. 3 shows a baseline PI-based control architecture for the heat pump operating in cooling mode, consisting of three decentralized feedback loops. The first loop uses CF to regulate T_r to the setpoint r_1 using a PI compensator. The second loop uses the EEV to regulate T_d to a target value that is computed by a dynamic schedule based on CF and T_a. This feedback indirectly regulates T_{SH} to the desired constant, as described in [1], [17], and has advantages over conventional T_{SH} feedback, such as improved start-up sequencing and operation when $T_{SH} < 0$, which cannot be measured since $T_s = T_e$ in this condition. The third feedback loop regulates T_c to the setpoint r_3 using OFS. This feedback loop is active under certain operating conditions, e.g. when T_a is below a threshold. In other conditions the feedback loop is open and OFS is set to a constant value, subject to override logic. Each

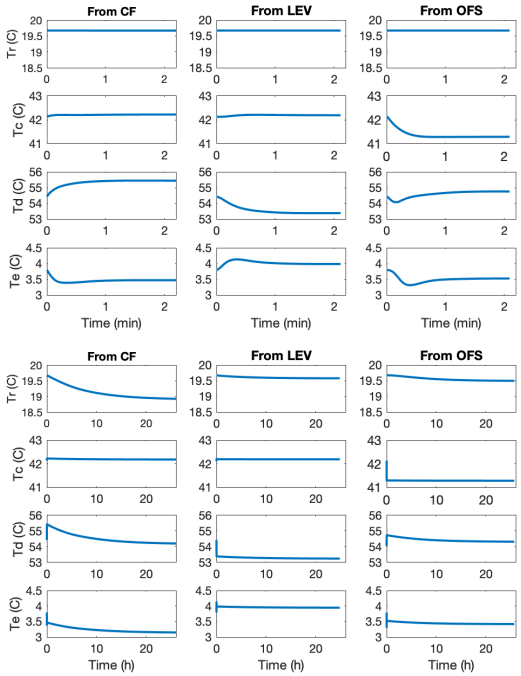

Fig. 2. Open-loop step responses on a 2 minute time scale (top 12) and a 24 hour time scale (bottom 12).

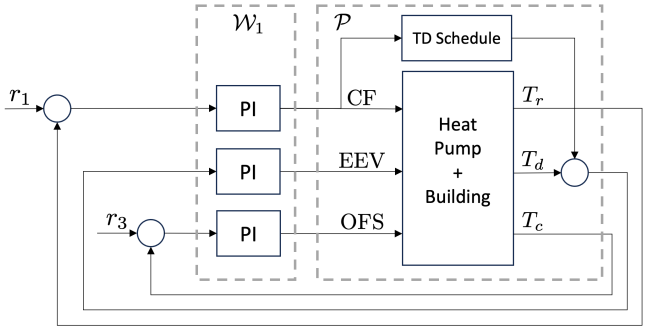

Fig. 3. PI control architecture for the heat pump.

PI includes anti-windup logic to enforce actuator constraints. Selector logic is used to enforce the output constraints (d) by switching one or more loops to another PI compensator for each constraint, but details are omitted.

PI compensators were designed for each loop, one at a time, using classic loop-shaping methods to tune a total of 12 PI parameters. Disk margins [18] were computed for each loop and multivariable disk margins were also computed for 16 possible loop closures. This was done because the system may operate with any combination of feedback loops open or closed, during start-up or shut-down, or because of actuator saturation or a constrained output is at its limit.

IV. H-Infinity MPC

An $\mathcal{H}_\infty$ loop-shaping MPC is essentially an $\mathcal{H}_\infty$ loop-shaping compensator (for robustness) [19], [20], [21] realized as a real-time constrained optimization computation (for constraint enforcement) [5]. The methodology uses the frequency response of $y_1 = \mathcal{P}_u u + \mathcal{P}_d d$ where $\mathcal{P}_u$ represents the transfer function from inputs $u = [\text{CF, EEV, OFS}]^T$ to

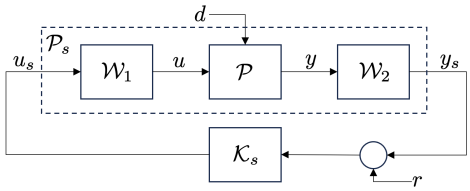

Fig. 4. $\mathcal{H}_\infty$ loop-shaped control architecture.

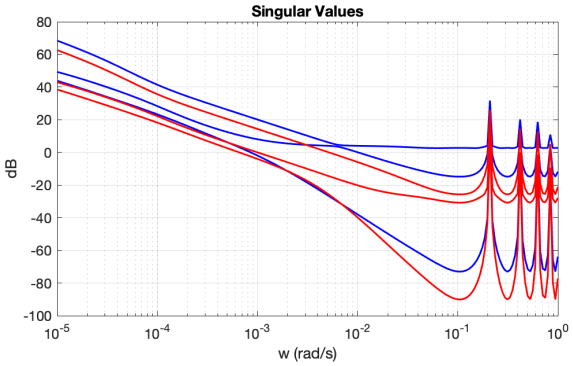

Fig. 5. Frequency response (singular values) for the shaped plant $\mathcal{W}_2\mathcal{P}_u\mathcal{W}_1$ (blue) (the PI loop gain) and the robustified loop gain $\mathcal{K}_s\mathcal{W}_2\mathcal{P}_u\mathcal{W}_1$ (red).

outputs $y_1 = [T_r, T_d, T_c]^T$, $\mathcal{P}_d$ is the disturbance transfer function from $d = [T_a, q_s, q_l]^T$ to y_1, and $y = [y_1, y_2]^T$ where $y_2 = [T_s, T_e, \phi_r]^T$. The frequency response of $\mathcal{P}_u$ is *shaped* with an input weight (filter) $\mathcal{W}_1$ and an output weight (filter) $\mathcal{W}_2$, such that the compensated open-loop $\mathcal{P}_s = \mathcal{W}_2\mathcal{P}_u\mathcal{W}_1$ has characteristics that ensure the closed-loop system satisfies Requirements (a)-(b). The shaping process is done in lieu of conventional plant model augmentations used in offset-free MPC.

The PI compensators described in Section III were lumped together to form $\mathcal{W}_1$, as shown in Fig. 3, and $\mathcal{W}_2 = \mathcal{I}$, the identity. Then the robustifying compensator $\mathcal{K}_s$ was computed by solving two decoupled Riccati equations [21], giving the $\mathcal{H}_\infty$ loop-shaped feedback architecture shown in Fig. 4, and frequency response shown in Fig. 5. The shaped plant $[\mathcal{W}_2\mathcal{P}_u\mathcal{W}_1, \mathcal{P}_d]^T$ was realized in discrete time with sampling period $h = 30$s as

$$x_s(k+1) = A_s x_s(k) + B_s u_s(k) + B_{ds} d(k) \tag{1a}$$
$$y_s(k) = C_s x_s(k) \tag{1b}$$
$$v_s(k) = F_s x_s(k)\,, \tag{1c}$$

where the output v_s includes all of the constrained variables (since (1) includes $\mathcal{W}_1$, whose constrained outputs are u). The $\mathcal{H}_\infty$ loop-shaping compensator $\mathcal{K}_s$ has an observer-based state feedback structure [21], realized as

$$\widehat{x}_s(k+1) = A_s\widehat{x}_s(k) + B_s u_s(k) + H_s\left(\widehat{y}_s(k) - y_s(k) + r(k)\right) \tag{2a}$$
$$\widehat{y}_s(k) = C_s\widehat{x}_s(k) \tag{2b}$$
$$u_s(k) = K_s\widehat{x}_s(k)\,, \tag{2c}$$

where $r = [r_1, 0, r_3]^T$, (2a)-(2b) are the state estimator equations with H_s being the robust estimator gain, and (2c)

is the state feedback.

In $\mathcal{H}_\infty$ loop-shaped MPC, we compute $Q_s \geq 0$, $R_s > 0$ and $P_s > 0$ for the quadratic cost function

$$J(u_s) = \sum_{k=1}^{N} x_s^T(k) Q_s x_s(k) + u_s^T(k) R_s u_s(k) + x_s^T(N) P_s x_s(N) \tag{3}$$

that is inverse optimal for the state feedback (2c), and then use (3) together with the output constraints $v_s \in \mathcal{V}$, where $\mathcal{V}$ is a set suitably defined to represent requirements (c) and (d), to numerically compute u_s as a constrained QP problem over the N-step receding horizon, where P_s satisfies

$$A_s^T P_s A_s - A_s^T P_s B_s (R_s + B_s^T P_s B_s)^{-1} B_s^T P_s A_s + Q_s = P_s \tag{4}$$

and

$$K_s = -(B_s^T P_s B_s + R_s)^{-1} B_s^T P_s A_s. \tag{5}$$

The state estimator (2a)-(2b) is used to initialize the MPC at each sample. See [5] for details.

In this realization, for constant values of disturbance d and reference r, the controller state $\widehat{x}_s$ converges to zero. This is a consequence of the stability of the closed-loop and observability of $\mathcal{K}_s$, and implies that a prediction model based upon only (1) using initialization from (2) cannot predict the constrained outputs v_s. Therefore, the prediction model was constructed using (1) augmented with a model that separately estimates the constrained outputs v_s subject to the unmeasured d, as follows.

With the original 18^{th}-order plant $\mathcal{P}$ realized as

$$x(k+1) = Ax(k) + B_u u(k) + B_d d(k) \tag{6a}$$
$$y(k) = Cx(k) \tag{6b}$$
$$v_s(k) = Fx(k) + Gu(k) \tag{6c}$$

augment the state x to include d in the usual way,

$$\bar{x}(k+1) = A_q \bar{x}(k) + B_q u(k) \tag{7a}$$
$$y(k) = C_q \bar{x}(k) \tag{7b}$$
$$v_s(k) = F_q \bar{x}(k) + Gu(k) \tag{7c}$$

where

$$A_q = \begin{bmatrix} A & B_d \\ 0 & I \end{bmatrix} \qquad B_q = \begin{bmatrix} B_u \\ 0 \end{bmatrix}, \tag{8}$$

$C_q = [C \;\; 0]$, $F_q = [F \;\; 0]$, and $\bar{x} = [x \;\; d]^T$, and construct a Luenberger observer, after verifying (A_q, C_q) is detectable,

$$\widehat{\bar{x}}(k+1) = A_q \widehat{\bar{x}}(k) + B_q u(k) + L_q \left(\widehat{y}(k) - y(k)\right) \tag{9a}$$
$$\widehat{y}(k) = C_q \widehat{\bar{x}}(k) \tag{9b}$$
$$\widehat{v}_s(k) = F_q \widehat{\bar{x}}(k) + Gu(k), \tag{9c}$$

with $Q_q \geq 0$ and $R_q > 0$. This was used in the prediction model to estimate the constraints v_s over the horizon. At each step k, the estimators (2a) and (9) were used to initialize the prediction model, which was used to evaluate cost (3) and constraints v_s (9c) over the N step horizon. The constrained QP was numerically solved using an ADMM algorithm [22] to compute u_s over the horizon to minimize (3) subject to $v_s \in \mathcal{V}$.

The design procedure is relatively simple: Once the PI compensators are designed (6 parameters), the only parameters to be tuned are Q_q and R_q for the constraint estimator (9), the horizon length N, and a softening parameter σ associated with the QP solver for the three true output constraints in v_s (the three true input constraints are hard and always feasible since the PI is invertible). All other parameters (K_s, H_s, P_s, Q_s, R_s, L_q) are computed automatically.

V. Offset-Free MPC

In this section, we use the linear time-invariant (LTI) model (6) obtained by linearization of the same Modelica plant, but adopt the standard offset-free MPC notation

$$x^+ = Ax + Bu + w \qquad y = Cx + \nu$$

where $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$, and $y \in \mathbb{R}^p$, with process and measurement disturbances $w \in \mathbb{R}^n$ and $\nu \in \mathbb{R}^p$, respectively. Unlike Section IV, offset-free MPC augments the model with integrating disturbance states for offset-free tracking; here, d denotes the augmented disturbance state, not the exogenous disturbance input.

Standard linear MPC on a nominal model exhibits steady-state offset under nonzero disturbances and plant–model error. Motivated by the Internal Model Principle, offset-free MPC augments the model with integrating disturbance states $d \in \mathbb{R}^{n_d}$ to obtain the following augmented dynamics and output equations

$$\begin{bmatrix} x \\ d \end{bmatrix}^+ = \underbrace{\begin{bmatrix} A & B_d \\ 0 & I \end{bmatrix}}_{:=\hat{A}_d} \begin{bmatrix} x \\ d \end{bmatrix} + \underbrace{\begin{bmatrix} B \\ 0 \end{bmatrix}}_{:=\hat{B}_d} u + \hat{w}$$
$$y = \underbrace{\begin{bmatrix} C & C_d \end{bmatrix}}_{:=\hat{C}_d} \begin{bmatrix} x \\ d \end{bmatrix} + \hat{\nu}$$

with design matrices $B_d \in \mathbb{R}^{n \times n_d}$, $C_d \in \mathbb{R}^{p \times n_d}$ and augmented disturbances $\hat{w} \sim \mathcal{N}(0, Q_{wd})$, $\hat{\nu} \sim \mathcal{N}(0, R_v)$, where Q_{wd} is the joint state–disturbance covariance and R_v the measurement noise covariance [23], [24]. The offset-free MPC controller then consists of three components as shown in Fig. 6:

a) State and disturbance estimator: The estimated states and disturbances $\hat{x}_d := [\hat{x}, \hat{d}]^{\mathrm{T}}$ are obtained using a stable estimator L_d (e.g., a Kalman filter):

$$\hat{x}_d^+ = \hat{A}_d \hat{x}_d + \hat{B}_d u + L_d \Big(y - \hat{C}_d \hat{x}_d \Big)$$

In this study, we design L_d as a discrete-time Kalman filter with process and measurement noise covariances Q_{wd} and R_v, respectively.

b) Target selector: The steady-state problem assumes a constant disturbance, $\hat{d}_{\mathrm{t}} = \hat{d}(k)$, and selects steady-state targets $(x_{\mathrm{t}}, u_{\mathrm{t}})$ with slacks $s_{\mathrm{t}} = [s_{r+}^\top, s_{r-}^\top, s_f^\top]^\top$ so that the

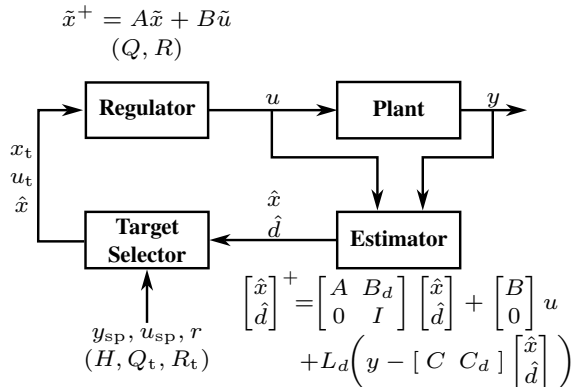

Fig. 6. Offset-free MPC architecture showing the estimator, steady-state target selector, and regulator. Adapted from Fig. 1.6 of [25].

regulated outputs $z = Hy \in \mathbb{R}^{n_r}$ track the reference r in the presence of disturbance (if possible):

$$\begin{aligned}
\min_{x_\mathrm{t},u_\mathrm{t},s_\mathrm{t}} \; & \frac{1}{2}\Big(|u_\mathrm{t} - u_\mathrm{sp}|^2_{R_\mathrm{t}} + |\hat{y}_\mathrm{t} - y_\mathrm{sp}|^2_{Q_\mathrm{t}} + |s_\mathrm{t}|^2_{S_\mathrm{t}}\Big) + c_\mathrm{t}^\top s_\mathrm{t} \\
\text{s.t. } & (I - A)x_\mathrm{t} - Bu_\mathrm{t} = B_d\hat{d}_\mathrm{t} \\
& \pm H\hat{y}_\mathrm{t} \le \pm r + s_{r\pm} \\
& Eu_\mathrm{t} \le e \\
& F\hat{y}_\mathrm{t} \le f - f_b + s_f \\
& \hat{y}_\mathrm{t} = Cx_\mathrm{t} + C_d\hat{d}_\mathrm{t} \\
& s_\mathrm{t} \ge 0
\end{aligned}$$

H selects offset-free outputs; (E, e) and (F, f) define input and output constraints; f_b is a backoff ensuring regulator stability at steady state [8]; $Q_\mathrm{t} \succeq 0$ and $R_\mathrm{t} \succ 0$ are penalties on non-offset-free outputs and inputs, respectively; $S_\mathrm{t} \succeq 0$ and $c_\mathrm{t} > 0$ encode quadratic and linear slack penalties, respectively.

c) Regulator: Given steady-state targets $(x_\mathrm{t}, u_\mathrm{t})$, the regulation MPC is posed in deviation variables with initial condition $\tilde{\boldsymbol{x}}(0) = \hat{x}(j) - x_\mathrm{t}$ over a horizon N, j being the current time index:

$$\begin{aligned}
\min_{\tilde{\boldsymbol{u}},\tilde{\boldsymbol{s}}} \; & \sum_{k=0}^{N-1} \ell(\tilde{\boldsymbol{x}}(k), \tilde{\boldsymbol{u}}(k), \tilde{\boldsymbol{s}}(k)) + \ell_f(\tilde{x}(N), \tilde{s}(N)) \\
\text{s.t. } & \tilde{x}^+ = A\tilde{x} + B\tilde{u} \\
& E\tilde{u} \le e - Eu_\mathrm{t} \\
& FC\tilde{x} \le f - F(Cx_\mathrm{t} + C_d\hat{d}_\mathrm{t}) + \tilde{s} \\
& \tilde{s} \ge 0 \\
& \ell(\tilde{x}, \tilde{u}, \tilde{s}) = \frac{1}{2}\Big(|\tilde{x}|^2_Q + |\tilde{u}|^2_R + |\tilde{s}|^2_S\Big) + c_{\tilde{s}}^\top \tilde{s} \\
& \ell_f(\tilde{x}, \tilde{s}) = \frac{1}{2}\Big(|\tilde{x}|^2_{Q_f} + |\tilde{s}|^2_S\Big) + c_{\tilde{s}}^\top \tilde{s}
\end{aligned}$$

$Q, Q_f \succeq 0$ are state and terminal penalty matrices from the target x_t; $R \succ 0$ is the input penalty from the target u_t; $S \succeq 0$ and $c_{\tilde{s}} > 0$ are quadratic and linear slack penalties for output constraints, respectively. The applied input is $u(j) = u_\mathrm{t} + \tilde{\boldsymbol{u}}^0(0)$ where $\tilde{\boldsymbol{u}}^0$ is the optimized input sequence. For compactness, we show the regulator in its basic form; in the actual implementation, a standard augmented-state formulation is used to penalize the input rate-of-change Δu with weight R_Δ. Both the steady-state target and regulator problems are solved using MPCTools, CasADi and qpOASES [26], [27], [28].

Remark 1 (Detectability and offset-free guarantee.). The design matrices B_d, C_d can be chosen arbitrarily with no effect on controller performance so long as the augmented system is detectable and the estimator covariances are tuned appropriately [29]. Assuming (A, C) is detectable, a convenient test is

$$\operatorname{rank}\begin{bmatrix} I - A & -B_d \\ C & C_d \end{bmatrix} = n + n_d \tag{10}$$

which is equivalent to detectability of $(\hat{A}_d, \hat{C}_d)$ [30]. Finally, when $n_d = p$ and Eq. (10) holds, the offset-free MPC controller has the standard property: if the closed loop is asymptotically stable and no constraints are active at steady state, the regulated outputs converge to their references without offset [30].

$$\lim_{k\to\infty} y(k) = y_\mathrm{t}, \qquad Hy_\mathrm{t} = r$$

A. LTI controller representation for disk margin analysis

Robustness assessment with disk margins relies on an explicit plant–controller interconnection. Disk margins quantify gain and phase uncertainty at the plant input–output interface, and their computation assumes a linear time-invariant interconnection. Thus, we derive the LTI representation of the offset-free MPC law in the unconstrained case.

Lemma 2 (Unconstrained LTI controller representation)**.** *In the absence of active constraints, offset-free MPC reduces to a LTI dynamic controller. Its state is the augmented current filter prior estimate* $\hat{x}_{d-} = \begin{bmatrix}\hat{x}_-\\ \hat{d}_-\end{bmatrix}$ *and the dynamics are*

$$\hat{x}^+_{d-} = A_c\hat{x}_{d-} + B_c y \qquad u = -\big(C_c\hat{x}_{d-} + D_c y\big) \tag{11}$$

with controller matrices

$$A_c = (\hat{A}_d - \hat{B}_d K_d)(I - L_d\hat{C}_d) \tag{12}$$
$$B_c = (\hat{A}_d - \hat{B}_d K_d)L_d \tag{13}$$
$$C_c = K_d(I - L_d\hat{C}_d) \tag{14}$$
$$D_c = K_d L_d \tag{15}$$

where K_d *is the modified control gain defined in Section VIII.*

The LTI representation is derived in Section VIII.

B. Tuning Parameters

Estimator: joint state-disturbance covariances Q_{wd}, and measurement noise covariance R_v.

Target selector: Steady-state penalty matrices Q_t, R_t, linear/quadratic slack penalties c_t, S_t, and backoffs f_b.

Regulator: Target-tracking penalty matrices Q, R, terminal weight Q_f, rate-penalty matrix R_Δ, linear/quadratic slack penalties S, $c_{\tilde{s}}$, horizon N, and sampling period h.

Remark 3 (Ease of tuning). The regulator and target selector weights (Q, R, R_Δ, Q_t, R_t) and slack penalties

are straightforward to adjust according to control priorities, desired behavior, and acceptable disk margins. In particular, increasing the rate penalty R_Δ reduces control aggressiveness and, in our study, was especially effective in improving robustness as measured by disk margins, at the expense of slower control action. We set the terminal penalty Q_f to the infinite-horizon linear quadratic regulator (LQR) cost-to-go matrix and used a horizon of $N = 5$ with a sampling period $h = 30$s. The estimator covariances were tuned in open-loop step tests to obtain sufficiently fast, smooth estimator dynamics, analogous to classical observer tuning [2]. Alternative approaches in the literature use process data or robustness criteria to guide covariance selection [9], [31].

Remark 4 (Linear slack weights and priority). We choose linear slack weights larger than any expected Lagrange multiplier for a given constraint to enforce satisfaction whenever feasible. To encode priorities, we scale the weights according to the relative importance of the constraints, assigning larger values to higher-priority equalities and smaller values to lower-priority ones. For this study, we prioritize output inequality constraints (equally weighted) over the offset-free equality constraints. This weighting ensures the controller satisfies output constraints first, then tracks as many offset-free setpoints as possible.

VI. Comparison

A. Setpoint Tracking and Disturbance Rejection

The controller design, estimator tuning, robustness analysis, and closed-loop simulations in this work are all based on a linearization about a nominal operating condition. Thus, the results should be interpreted as local to operating regimes for which this approximation remains representative. We compare the three controllers – the baseline PI, $\mathcal{H}_\infty$ MPC, and offset-free MPC – with the simulation test shown in Fig. 7 featuring both a setpoint change in the room temperature and a sensible heat disturbance. Input limits are as follows: compressor frequency (CF) 20–120 Hz, expansion valve (EEV) 50–500 counts, and outdoor fan speed (OFS) 100–750 rpm. Output constraints include the evaporator temperature limit $T_e > 0°\mathrm{C}$, discharge temperature limit $T_d < 100°\mathrm{C}$, and discharge superheat limit $T_{DSH} > 10°\mathrm{C}$.

a) Setpoint change: At the start of the test, the room temperature setpoint is stepped from $20°\mathrm{C}$ to $18°\mathrm{C}$. All three controllers have pull-down times less than one hour, and all three controllers transiently drive T_e to the $0°\mathrm{C}$ lower limit before settling at a steady state temperature, which rate limits the pull-down time. Regarding the inputs, all three controllers show similar input profiles in the compressor frequency and expansion valve. However the outdoor fan speed behavior differs. While PI and $\mathcal{H}_\infty$ maintain a high fan speed, the offset-free MPC lowers the fan speed to its minimum. This is counterintuitive because it would seemingly reduce heat transfer through the heat pump. However, when at the evaporator temperature limit, lowering the outdoor fan speed reduces the refrigerant subcooling, which reduces the amount of refrigerant in the outdoor unit, which therefore is transferred to the indoor unit, so more of the evaporator is in the two-phase state, actually increasing heat transfer. From a controls perspective, the fan speed reduction raises the condensing temperature, which creates slack in the $T_e \geq 0°\mathrm{C}$ constraint; offset-free MPC then spends that slack by increasing the compressor frequency to pump more heat out of the room. After the setpoint transition is largely completed, the fan speed increases again as the controller settles to the new operating condition.

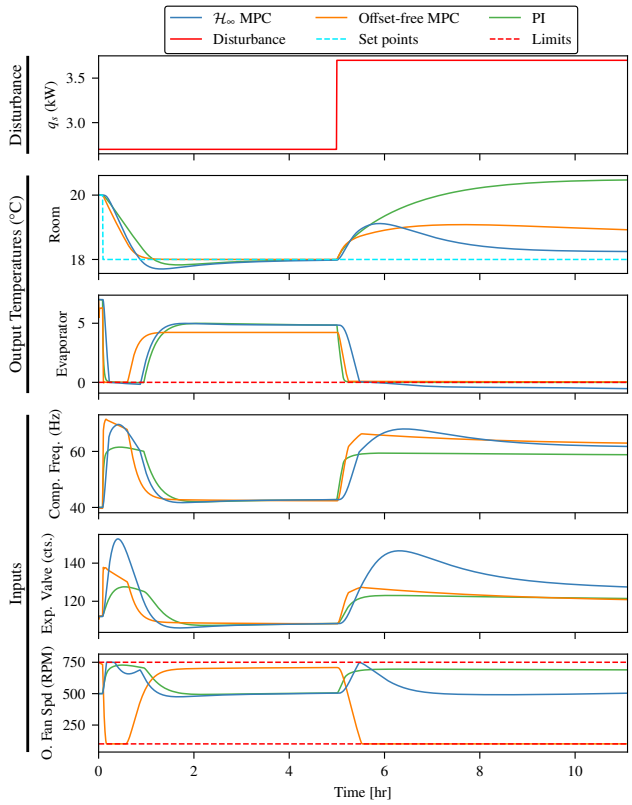

Fig. 7. Setpoint tracking and disturbance rejection for the three controllers.

b) Disturbance rejection: A $+1\,\mathrm{kW}$ sensible heat-load step is applied at the new steady state. Due to the lower limit on the evaporator temperature for all three controllers, the room temperature departs from the $18°\mathrm{C}$ setpoint and does not recover within the simulation time. Among the controllers, $\mathcal{H}_\infty$ MPC exhibits the smallest sustained offset, offset-free MPC shows a similar excursion with slower decay, and PI drifts farthest from the setpoint. However, studying the evaporator temperature reveals that while offset-free MPC and PI maintain the evaporator at the $0°\mathrm{C}$ constraint, $\mathcal{H}_\infty$ MPC has a slight persistent constraint violation due to the soft constraint enforcement. Similar to the setpoint change inputs, the compressor frequency and expansion valves share similar profiles. Again, while the $\mathcal{H}_\infty$ MPC and PI controllers maintain medium to high outdoor fan speeds, offset-free MPC drops the fan speed to its minimum, creating evaporator slack that permits higher compressor frequency.

In summary, offset-free MPC and $\mathcal{H}_\infty$ MPC achieve similar pull-down times of around 45 minutes compared to 1 hour for PI, while the integrated absolute room-temperature

TABLE II

DISK MARGINS.

Control	OFS on	OFS off
PI	1.35	1.51
$\mathcal{H}_\infty$ MPC	1.23	1.07
Offset-free MPC	1.42	1.42

TABLE III

TUNABLE PARAMETER COUNT.

Control	Tuning Parameters	Number*
PI	$6 \times (k_p, k_i)$	12
$\mathcal{H}_\infty$ MPC	$3 \times (k_p, k_i), Q_q, R_q, \sigma, N$	10
Offset-free MPC	$Q_{wd},\ R_v,\ Q_{\mathrm{t}},\ R_{\mathrm{t}},\ S_{\mathrm{t}},\ c_{\mathrm{t}},\ f_b,$ $Q,\ Q_f,\ R,\ R_\Delta,\ S,\ c_{\tilde{s}},\ N$	64
*Assuming a structure to the matrices		

tracking errors are 13.5, 4.6, and 6.7 °C·hr for PI, $\mathcal{H}_\infty$ MPC, and offset-free MPC, respectively. However, only $\mathcal{H}_\infty$ MPC violates the evaporator-temperature constraint, with T_e dropping to $-0.5°\mathrm{C}$.

B. Disk Margin Assessment

Disk margins quantify simultaneous gain and phase uncertainty tolerance by considering the minimum perturbation in input/output channels that destabilizes the closed-loop system [18]. For constrained MPC, this analysis is inherently local, since each active-constraint set defines a local controller; accordingly, the reported margins apply only while the active set remains unchanged. Moreover, these margins are evaluated for the linearized plant model used in the controller design. The multiloop disk margin can be computed using the structured singular value.

Given transfer functions for the plant $P(s)$ and controller $K(s)$, the input and output loop transfer functions are

$$L_{\mathrm{i}}(s) = K(s)P(s), \qquad L_{\mathrm{o}}(s) = P(s)K(s)$$

For MIMO systems, the sensitivity is $S(s) = (I + L(s))^{-1}$ defined with respect to either loop $L(s)$. The multivariable disk margin extends the SISO notion of simultaneous gain and phase margin: it is the largest disk in the complex plane centered at unity such that $(I + L(s)D)^{-1}$ remains stable for all perturbations D in the disk. To compute this margin, introduce a skew parameter $\sigma \in \mathbb{R}$ that shifts the disk center. The associated skewed sensitivity is

$$M(s) = S(s) + \tfrac{\sigma-1}{2} I$$

where $\sigma > 0$ favors gain increase, $\sigma < 0$ favors gain decrease, and $\sigma = 0$ gives the balanced case. With structured diagonal uncertainty set

$$\mathbf{\Delta} \coloneqq \mathrm{diag}(\delta_1 I_{n_1}, \ldots, \delta_m I_{n_m}) : \delta_i \in \mathbb{C}, |\delta_i| \le \alpha,\ \alpha > 0$$

the disk margin can be written in terms of the structured singular value

$$\mu_\Delta(M) = \frac{1}{\min_{\Delta \in \mathbf{\Delta}} \{|\Delta| : \det(I - M\Delta) = 0\}}$$

which measures the smallest destabilizing perturbation consistent with the block structure. The multiloop disk margin is then

$$\alpha_{\max} = \left(\sup_\omega \mu_\Delta(M(j\omega)) \right)^{-1}$$

We focus on the balanced input case in what follows and compute upper bounds for the structured singular value [32].

All controllers comfortably exceed a required disk margin of ~ 0.7 as shown in Table II, which reports input disk margins for the balanced case when the outdoor fan speed is included in the feedback loop, and also when it is at a constant value and the OFS feedback loop is open. It is mildly surprising that the offset-free MPC controller has the largest nominal disk margin and the $\mathcal{H}_\infty$ MPC margin is smaller, but the $\mathcal{H}_\infty$ controller maximizes a different metric and the difference is not significant. Table III summarizes lower-bound counts of tunable parameters for the three architectures.

VII. CONCLUSIONS

This paper studied two alternative MPC architectures for control of a heat pump in terms of transient performance, robustness to plant uncertainty and design methodology. All three controllers meet robustness margin requirements (this plant is relatively forgiving), and more importantly provide for evaluation of disk margins that are important for evaluation of industrial controllers. The selector-based PI architecture possesses the fewest parameters to tune, and all may be designed via classical loop shaping provided a model is available. Yet the architecture itself is inflexible and needs to be modified for different types of products and modes of operation, for which there is significant engineering expense. The $\mathcal{H}_\infty$ loop-shaped MPC seemingly offers the same design methods and number of parameters to tune, although there are additional parameters related to soft constraints in the optimizer that affect performance, and the disturbance observer also requires tuning of parameters. The offset-free MPC achieves interesting, non-intuitive performance as seen by its adjustment of outdoor fan speed during a transient simulation, which is a clear advantage of multivariable MPC in general and certainly in this particular case. It also achieves similar robustness margins to the other two designs, although this is accomplished by expert tuning of parameters in the controller and estimator. But with increasingly available data measured from heat pumps operating in the field, automated covariance tuning [9] offers promise, and future work will focus on nonlinear simulation testing and automated covariance tuning.

REFERENCES

[1] S. A. Bortoff and K. Tsuji, "Robust unfalsified control of a heat pump," in *IEEE Conference on Control Technologies and Applications*, 2024.

[2] M. Wallace, B. Das, P. Mhaskar, J. House, and T. Salsbury, "Offset-free model predictive control of a vapor compression cycle," *Journal of Process Control*, vol. 22, no. 7, pp. 1374–1386, 2012.

[3] D. Leducq, J. Guilpart, and G. Trystram, "Non-linear predictive control of a vapour compression cycle," *International Journal of Refrigeration*, vol. 29, no. 5, pp. 761–772, 2006.

[4] M. S. Elliott and B. P. Rasmussen, "Decentralized model predictive control of a multi-evaporator air conditioning system," *Control Engineering Practice*, vol. 21, pp. 1665–1677, 2013.

[5] S. A. Bortoff, P. Schwerdtner, C. Danielson, S. D. Cairano, and D. J. Burns, "H-Infinity Loop-Shaped Model Predictive Control with HVAC Application," *IEEE Transactions on Control Systems Technology*, vol. 30, no. 5, pp. 2188–2203, March 2022.
[6] K. R. Muske and T. A. Badgwell, "Disturbance modeling for offset-free linear model predictive control," *Journal of Process Control*, vol. 12, no. 5, pp. 617–632, 2002.
[7] U. Maeder, F. Borrelli, and M. Morari, "Linear offset-free model predictive control," *Automatica*, vol. 45, pp. 2214–2222, 2009.
[8] S. J. Kuntz and J. B. Rawlings, "Offset-free model predictive control: stability under plant-model mismatch," 2024, arXiv:2412.08104 [eess].
[9] ——, "Maximum Likelihood Identification of Linear Models with Integrating Disturbances for Offset-Free Control," *IEEE Transactions on Automatic Control*, pp. 1–14, 2025.
[10] *Modelica Language Specification Version 3.5*, Modelica Association, https://www.Modelica.org/, Feb. 2021.
[11] P. Fritzon, *Principles of Object Oriented Modeling and Simulation with Modelica 3.3: A Cyber-Physical Approach*. Wiley, 2015.
[12] S. A. Bortoff, H. Qiao, and C. R. Laughman, "Modeling and control of a multi-mode heat pump," in *Proceedings of the 8th IEEE Conference on Control Technology and Application*, 2024.
[13] H. Qiao, V. Aute, and R. Radermacher, "Transient modeling of a flash tank vapor injection heat pump system–Part I: Model development," *International Journal of Refrigeration*, vol. 49, pp. 169–182, 2015.
[14] H. Qiao, C. R. Laughman, S. A. Bortoff, and D. J. Burns, "Dynamic characteristics of an R410A multi-split variable refrigerant flow air conditioning system," in *Proceedings of the 12th IEA Heat Pump Conference*, 2017.
[15] M. Wetter, W. Zuo, T. S. Nouidui, and X. Pang, "Modelica buildings library," *Journal of Building Performance Simulation*, vol. 7, no. 4, pp. 253–270, 2014.
[16] S. Zhan, A. Chakrabarty, C. Laughman, and A. Chong, "A virtual testbed for robust and reproducible calibration of building energy simulation models," in *18th IBPSA International Conference and Exhibition Building Simulation*, Sept. 2023.
[17] D. J. Burns, C. R. Laughman, and S. Bortoff, "System and method for controlling vapor compression systems," US Patent 10,495,364, Dec. 3 2019.
[18] P. Seiler, A. Packard, and P. Gahinet, "An introduction to disk margins," *IEEE Control Systems Magazine*, pp. 78–95, Oct. 2020.
[19] K. Glover and D. McFarlane, "Robust stabilization of normalized coprime factor plant descriptions with H_∞ bounded uncertainty," *IEEE Transactions on Automatic Control*, vol. 34, no. 8, pp. 821–830, 1989.
[20] G. Vinnicombe, *Uncertainty and Feedback: H_∞ Loop-Shaping and the ν-Gap Metric*. Imperial College Press, 2001.
[21] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*. Wiley, 2005.
[22] A. Raghunathan and S. D. Cairano, "Infeasibility detection in alternative direction method of multipliers for convex quadratic programs," in *IEEE Conference on Decision and Control*, 2014.
[23] E. J. Davison and H. W. Smith, "Pole assignment in linear time-invariant multivariable systems with constant disturbances," *Automatica*, vol. 7, pp. 489–498, 1971.
[24] H. Kwakernaak and R. Sivan, *Linear Optimal Control Systems*. New York: John Wiley and Sons, 1972.
[25] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd ed. Santa Barbara, CA: Nob Hill Publishing, 2020, 770 pages, ISBN 978-0-9759377-5-4.
[26] M. J. Risbeck and J. B. Rawlings, "MPCTools: Nonlinear model predictive control tools for CasADi (Python inferface)," 2015. [Online]. Available: https://bitbucket.org/rawlings-group/mpc-tools-casadi
[27] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi—a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, March 2019.
[28] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, "qpOASES: a parametric active-set algorithm for quadratic programming," *Mathematical Programming Computation*, vol. 6, no. 4, pp. 327–363, 2014.
[29] M. R. Rajamani, J. B. Rawlings, and S. J. Qin, "Achieving state estimation equivalence for misassigned disturbances in offset-free model predictive control," *AIChE Journal*, vol. 55, no. 2, pp. 396–407, February 2009.
[30] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free MPC control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.
[31] G. Pannocchia and A. Bemporad, "Combined design of disturbance model and observer for offset-free model predictive control," *IEEE Transactions on Automatic Control*, vol. 52, no. 6, pp. 1048–1053, June 2007.
[32] M. K. Fan, A. L. Tits, and J. C. Doyle, "Robustness in the Presence of Joint parametric Uncertainty and Unmodeled Dynamics," in *1988 American Control Conference*, 1988, pp. 1195–1200.
[33] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. Englewood Cliffs, N. J.: Prentice-Hall, 1979.

VIII. Proof of Lemma 2

Proof: The LTI realization in Eqs. (11) to (15) follows from the standard current-filter observer form, in which the control law is first written as $u = -K_d\hat{x}_d$ in terms of the posterior estimate $\hat{x}_d$, and then realized with prior state $\hat{x}_{d-}$ [33]. It remains to derive the offset-free MPC gain K_d.

A. Steady-State Target Optimization

Without loss of generality, we take $u_{\mathrm{sp}}, y_{\mathrm{sp}}, r = 0$. For the unconstrained case, the steady-state targets solve:

$$\min_{x_{\mathrm{t}},u_{\mathrm{t}}} \frac{1}{2}\left(|u_{\mathrm{t}}|^2_{R_{\mathrm{t}}} + \left|Cx_{\mathrm{t}} + C_d\hat{d}_{\mathrm{t}}\right|^2_{Q_{\mathrm{t}}}\right)$$
$$\text{s.t.} \begin{bmatrix} I-A & -B \\ HC & 0 \end{bmatrix}\begin{bmatrix} x_{\mathrm{t}} \\ u_{\mathrm{t}} \end{bmatrix} = \begin{bmatrix} B_d\hat{d}_{\mathrm{t}} \\ -HC_d\hat{d}_{\mathrm{t}} \end{bmatrix}$$

Solving the equality-constrained least-squares problem, we obtain the optimal solution $x^*_{\mathrm{t}}, u^*_{\mathrm{t}}$ and the Lagrange multipliers λ^*_1, λ^*_2 as

$$\begin{bmatrix} x^*_{\mathrm{t}} \\ u^*_{\mathrm{t}} \\ \lambda^*_1 \\ \lambda^*_2 \end{bmatrix} = \begin{bmatrix} K_{dx} \\ K_{du} \\ K_{\lambda,1} \\ K_{\lambda,2} \end{bmatrix}\hat{d}_{\mathrm{t}}$$

$$\begin{bmatrix} K_{dx} \\ K_{du} \\ K_{\lambda,1} \\ K_{\lambda,2} \end{bmatrix} := \begin{bmatrix} C^TQ_{\mathrm{t}}C & 0 & (I-A)^T & (HC)^T \\ 0 & R_{\mathrm{t}} & (-B)^T & 0 \\ I-A & -B & 0 & 0 \\ HC & 0 & 0 & 0 \end{bmatrix}^{\dagger}\begin{bmatrix} -C^TQ_{\mathrm{t}}C_d \\ 0 \\ B_d \\ -HC_d \end{bmatrix}$$

† denotes the Moore-Penrose pseudoinverse.

Given steady-state targets, the offset-free control law is:

$$u = -K(\hat{x} - x^*_{\mathrm{t}}) + u^*_{\mathrm{t}} = -K\hat{x} + \begin{bmatrix} K & I \end{bmatrix}\begin{bmatrix} x^*_{\mathrm{t}} \\ u^*_{\mathrm{t}} \end{bmatrix}$$

Since the regulator problem is unconstrained, the control gain K is the unconstrained infinite-horizon LQR gain for the system (A, B) with cost weights Q, R. When an input-rate penalty is included, the steady-state target analysis remains unchanged, since $\Delta u = 0$ at steady state, while only the regulator gain is modified through the standard augmented system with state $\begin{bmatrix} x^\top & u_-^\top \end{bmatrix}^\top$ and manipulated variable Δu. Substituting the steady-state target expression, we obtain

$$u = -K_d\hat{x}_d$$

where $K_d = \begin{bmatrix} K & -(KK_{dx} + K_{du}) \end{bmatrix}$. ∎