H Infinity Loop-Shaped Model Predictive Control with Heat Pump Application

Bortoff, S.A.; Schwerdtner, P.; Danielson, C.; Di Cairano, S.

TR2019-052 June 29, 2019

Abstract

In this paper we derive a formulation for Model Predictive Control (MPC) of linear time-invariant systems based on H infinity loop-shaping. The design provides an optimized stability margin for problems that require state estimation. Input and output weights are designed in the frequency domain to satisfy steady-state and transient performance requirements, in lieu of conventional MPC plant model augmentations. The H infinity loop-shaping synthesis results in an observer-based state feedback structure. Using the linear state feedback law, an inverse optimal control problem is solved to design the MPC cost function, and the H infinity state estimator is used to initialize the prediction model at each time step. The MPC inherits the closed-loop performance and stability margin of the loopshaped design when constraints are inactive. We apply the methodology to a multi-zone heat pump system in simulation. The design rejects constant unmeasured disturbances and tracks constant references with zero steady-state error, has good transient performance, provides an excellent stability margin, and enforces input and output constraints.

European Control Conference (ECC)

$\mathcal{H}_\infty$ Loop-Shaped Model Predictive Control with Heat Pump Application

Scott A. Bortoff[1*], Paul Schwerdtner[2], Claus Danielson[1] and Stefano Di Cairano[1]

***Abstract*— In this paper we derive a formulation for Model Predictive Control (MPC) of linear time-invariant systems based on $\mathcal{H}_\infty$ loop-shaping. The design provides an optimized stability margin for problems that require state estimation. Input and output weights are designed in the frequency domain to satisfy steady-state and transient performance requirements, in lieu of conventional MPC plant model augmentations. The $\mathcal{H}_\infty$ loop-shaping synthesis results in an observer-based state feedback structure. Using the linear state feedback law, an inverse optimal control problem is solved to design the MPC cost function, and the $\mathcal{H}_\infty$ state estimator is used to initialize the prediction model at each time step. The MPC inherits the closed-loop performance and stability margin of the loop-shaped design when constraints are inactive. We apply the methodology to a multi-zone heat pump system in simulation. The design rejects constant unmeasured disturbances and tracks constant references with zero steady-state error, has good transient performance, provides an excellent stability margin, and enforces input and output constraints.**

I. INTRODUCTION

Model Predictive Control (MPC) is a feedback control methodology in which an optimization problem is solved in real-time to minimize a cost function, which represents performance, subject to a set of constraints, including the plant dynamics and limits on inputs, states and outputs. At each sample time, an optimizer computes a sequence of control inputs over a finite-time horizon that minimizes the cost function subject to the constraints. The first element of the control sequence is applied to the plant at that sample time, and the process is repeated at subsequent sample times in a recursive, receding horizon manner.

An MPC control law includes the following components:

1) A prediction model,

$$x_p(k+1) = A_p x_p(k) + B_p u(k) \tag{1a}$$
$$z(k) = E_p x_p(k) \tag{1b}$$
$$v(k) = F_p x_p(k) + G_p u(k)\,, \tag{1c}$$

where $x_p(k)$ is the generalized plant state, $u(k)$ is the control input, $z(k)$ is the performance output and $v(k)$ is the constrained output vector, at time kT, where T is the sample time. Here, *generalized* means that the plant model has been augmented in order to achieve performance specifications such as disturbance rejection or reference tracking. The prediction model computes the response of the generalized plant, specifically z and v, subject to a control input $[u(k), \ldots, u(k+N)]$ and to an initial condition $x_p(k)$, over the finite-time horizon of length N.

2) A cost function, typically quadratic,

$$J(u) = x_p^T(N) P x_p(N) + \sum_{k=1}^{N-1} z^T(k) Q z(k) + u^T(k) R u(k)\,, \tag{2}$$

where matrices $P > 0$, $Q \geq 0$ and $R > 0$ are designed to meet performance requirements.

3) Inequality constraints

$$v_{min} \leq v(k) \leq v_{max}\,, \tag{3}$$

which are derived from performance requirements and plant limitations.

4) An optimization algorithm that computes the cost-minimizing control sequence subject to the constraints at each time kT.

For many applications, such as the heat pump system we consider in this paper, a state estimator must be used to initialize the prediction model at each sample time, because the full state is not available from direct measurement. If the plant is subject to an unmeasured disturbance, then the state estimator is typically constructed by some kind of plant model augmentation, such as adding a constant offset vector to the output and estimating it together with the plant state to compensate for the effects of the unknown disturbance. It is well-known that this type of plant augmentation and estimator design will result in closed-loop offset-free tracking of the MPC, under minor assumptions [1], [2].

However, it is also well-known that using an estimated state in a full state feedback control law such as Linear Quadratic Gaussian (LQG), which is the basis of MPC because the cost (2) is quadratic and the estimator is a Kalman filter, provides no guarantee of any stability margin [3]. This fact has been generally ignored in the MPC literature, in which the typical design methodology emphasizes performance and constraint enforcement, but often overlooks robust stability, especially in the design of the plant augmentations, the state estimation gain, and the cost matrices Q and R.

This paper presents a design methodology that we call Loop-Shaped Model Predictive Control (LSMPC), in which the MPC design is based not on LQG but upon robust $\mathcal{H}_\infty$ loop-shaping [4], [5]. This provides an optimized stability margin, and provides a rigorous way to *design* the plant augmentations to meet performance requirements for closed-loop bandwidth, transient response, disturbance rejection and

[1]Mitsubishi Electric Research Laboratories, Cambridge, MA, 02139 USA (e-mail: {bortoff,danielson,dicairano}@merl.com)

[2]Technical University of Berlin, Straße des 17, Juni 135, 10623, Berlin, Germany, (e-mail: schwerdt@math.tu-berlin.de)

*Corresponding author

reference tracking. The methodology results in values for the state estimator gain, and for Q, R and P in (2). It consists of the following steps.

1) Design input and output weights for the plant to meet closed-loop performance specifications. This step *replaces* conventional MPC plant augmentations.
2) Synthesize the $\mathcal{H}_\infty$ loop-shaping compensator, which has an observer-based state feedback structure like LQG, but provides an optimized stability margin.[1]
3) Compute cost matrices $Q \geq 0$, $R > 0$ and $P > 0$ by solving an inverse-optimal control problem for the state feedback from step (2).
4) Use the $\mathcal{H}_\infty$ state estimator to initialize the prediction model at each time step. Some additional augmentations are necessary to enforce output constraints if any are present in the problem.

The closed-loop LSMPC inherits the disturbance rejection, reference tracking, transient response and stability margin of the $\mathcal{H}_\infty$ loop-shaped controller in the region where constraints are inactive, which normally, for good control designs, contains the target equilibria. The loop-shaped state estimator gain is computed by the $\mathcal{H}_\infty$ synthesis, and therefore does not require additional tuning. The key synthesis steps involve computing solutions to decoupled Riccati equations or solving a Linear Matrix Inequality (LMI), which are straightforward with available tools such as SeDuMi [6] and YALMIP [7]. Of course, LSMPC can also enforce constraints on inputs, states and outputs.

This paper extends several previously published results. Rowe and Maciejowski [8] apply $\mathcal{H}_\infty$ loop-shaping and inverse optimality to compute the state estimator and MPC cost. However, the authors do not consider output constraints, and they augment models of the reference and disturbance into the plant model, implying that the disturbance must be measured or estimated. In this paper, all plant augmentations are the result of $\mathcal{H}_\infty$ loop-shaping, we consider input and output constraints, and we consider unmeasured disturbances without the need to incorporate a disturbance model into the plant model. We remark that some conventional plant augmentations, such as adding a constant offset vector to the output and estimating its value in order to achieve offset-free tracking, make the augmented plant uncontrollable, which violates a sufficient condition for the $\mathcal{H}_\infty$ loop-shaping compensator. Maciejowski [9] and Di Cairano and Bemporad [10] consider the problem of controller matching, or finding a matching cost function, assuming an output feedback controller or state feedback controller, respectively, is given *a priori*. In this paper we *design* both the state feedback and state estimator to meet performance requirements and also to optimize a stability margin.

LSMPC may have some disadvantages. First, we have observed that for some plants it may result in a numerically ill-conditioned matrix P in (2), which can make the real-time optimization problem ill-conditioned, although there are numerical methods that can be employed to address or at least reduce this problem. And second, for some problems, the prediction model can have a larger dimension when compared to a conventional MPC design. This is because, in our formulation, any constrained outputs require an additional prediction model augmentation. However, in the special cases that the MPC need not enforce output constraints, or the problem does not require rejection of an unmeasured disturbance, the prediction model does not require this additional augmentation and could be of lower dimension than a conventional MPC design.

[1]Technically the resulting compensator possesses a slightly sub-optimal stability margin.

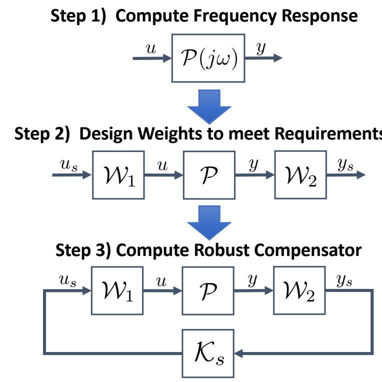

Fig. 1. Design steps for $\mathcal{H}_\infty$ loop-shaping. 1) Compute the frequency response of $\mathcal{P}$. 2) Design weights $\mathcal{W}_1$ and $\mathcal{W}_2$ to shape $|\mathcal{P}_s(j\omega)| = |\mathcal{W}_1(j\omega)\mathcal{P}(j\omega)\mathcal{W}_2(j\omega)|$. 3) Compute the robustifying compensator $\mathcal{K}_s$.

Throughout this paper, we use the calligraphic font to represent a linear system, e.g. $\mathcal{P}$, and the italics roman font to represent matrices, e.g., A_s. We employ the positive feedback convention that is standard in the $\mathcal{H}_\infty$ loop-shaping literature. In Section II we present the design steps for a general linear time-invariant system. In Section III we design the LSMPC for a multi-zone heat pump system which requires state estimation, is subject to an unmeasured disturbance, and has constraints on inputs and outputs. We suggest several extensions to the result in Section IV.

II. $\mathcal{H}_\infty$ Loop-Shaped MPC

Consider the scaled plant with minimal realization

$$x(k+1) = Ax(k) + Bu(k) + B_q q(k) \tag{4a}$$
$$y(k) = Cx(k) \tag{4b}$$
$$v(k) = Fx(k) + Gu(k), \tag{4c}$$

where $x(k) \in \mathbb{R}^n$, $u(k) \in \mathbb{R}^m$ is the control, $q(k) \in \mathbb{R}^d$ is an unmeasured disturbance, $y(k) \in \mathbb{R}^p$ is the measured output, $v(k) \in \mathbb{R}^l$ is the constrained output. We assume (4) is strictly proper. If not, then the $\mathcal{H}_\infty$ weights must be made strictly proper for the $\mathcal{H}_\infty$ synthesis.

To exemplify the design procedure, assume that the output vector $y \in \mathbb{R}^p$ is divided into two disjoint sets: A set of regulated outputs $y_1 \in \mathbb{R}^{p_1}$, where $p_1 \leq m$, and a set of other measurements $y_2 \in \mathbb{R}^{p-p_1}$, so that $y = [y_1^T \; y_2^T]^T$. The outputs y_1 are those which we wish to control, and would be incorporated into the performance output z in conventional

MPC, while the remainder of measurements are included in the feedback but are not considered "performance" variables.

Assume a typical set of closed-loop requirements:

A) Output y_1 must track constant reference r with zero steady-state error, if possible;
B) A constant disturbance q must be rejected with zero steady-state tracking error, if possible;
C) The closed-loop bandwidth must be ω_b for the regulated part of the system, if possible;
D) The closed-loop system must satisfy a set of transient response specifications (e.g. rise time), if possible;
E) Hard constraints must be enforced on the control input u; and
F) Constraints must be enforced on the outputs v.

In requirements A-D, by "if possible" we mean, whenever the constraints allow that to be achieved.

Following [4], [13], [14], $\mathcal{H}_\infty$ loop-shaping proceeds by computing the frequency response for the system $y = \mathcal{P}u$ (and also the frequency response for the system $y = \mathcal{P}_q q$), translating the bandwidth and transient specifications into the frequency domain, and designing an input weight $\mathcal{W}_1$ and an output weight $\mathcal{W}_2$, in order shape the frequency response of the compensated open-loop system $\mathcal{P}_s = \mathcal{W}_1\mathcal{P}\mathcal{W}_2$ to have characteristics that will ensure the nominal closed-loop system satisfies Requirements A-D. $\mathcal{K}_s$ is then computed by solving two decoupled Riccati equations. The basic steps are shown in Fig. 1.

It is common for $\mathcal{W}_1$ to include integral action to meet Requirements A and B, and for $\mathcal{W}_2$ to be a diagonal matrix of constants to account for scaling, although more complex weights can be used for decoupling for example. Less conventionally, integral action can instead be included in $\mathcal{W}_2$, and $\mathcal{W}_1$ can be a diagonal matrix of constants. This has some advantages for our MPC formulation. We describe each approach briefly below, referring the reader unfamiliar with the theory to the literature.

A. $\mathcal{W}_1$ Integral Action

Commonly, $\mathcal{W}_1$ is a diagonal system of Proportional-Integral (PI) type compensators, which we express in continuous-time Bode form as

$$\mathcal{W}_{1i}(s) = k_{1i}\frac{1 + s/\omega_{1i}}{s}\,, \tag{5}$$

for $1 \le i \le m$. The integral action ensures requirements A and B are satisfied, and the gain k_{1i} and zero location ω_{1i} are tuned iteratively using the frequency response (singular values) of $\mathcal{P}_s$ to achieve the specified cross-over frequency and a desirable cross-over phase in order to meet requirements C and D. (The continuous-time Bode parameters k_{1i} and ω_{1i} are the natural choice to achieve these specifications.) Higher-order weights or off-diagonal terms are possible but for simplicity of exposition, we assume (5) can meet the requirements. $\mathcal{W}_2$ is typically a diagonal system with constant gains k_{2i}, $1 \le i \le p_1$ corresponding to y_1. High-pass filters

$$\mathcal{W}_{2j}(s) = k_{2j}\frac{s}{1 + s/\omega_{2j}}\,, \tag{6}$$

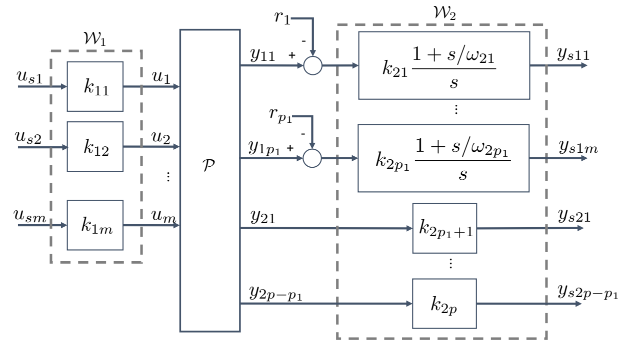

Fig. 2. The shaped plant $\mathcal{P}_s = \mathcal{W}_1\mathcal{P}\mathcal{W}_2$ with integral action in the $\mathcal{W}_2$ weight.

for $p_1 + 1 \le j \le p$ may be used for y_2. The zeros in (6) block DC, ensuring that y_1 tracks r with zero steady-state error, but allowing the high frequency content in y_2 to pass into the controller. This can improve the performance of the $\mathcal{H}_\infty$ estimator.

B. $\mathcal{W}_2$ Integral Action

An alternative is to incorporate integral action on the regulated outputs y_1. In this approach, $\mathcal{W}_1$ is a diagonal system of constants to compensate for any scaling issues, i.e., to allow for more or less weight on particular actuators as needed. $\mathcal{W}_2$ is a diagonal system in which the weights corresponding to y_1 are PI-type compensators (5), and weights corresponding to y_2 are constants or low-pass filters. In this case, reference r is injected at the output y_1, so that the PI weights act on $(y_1 - r)$, as shown in Fig. 2. This architecture is less common in the $\mathcal{H}_\infty$ literature, but it simplifies our MPC design, as will be explained below.

C. Comparison to Conventional Augmentations

Many conventional MPC augmentations, such as transforming control inputs to incremental form or adding integral action to output signals, are meant to meet steady-state tracking requirements or enforce constraints on rates-of-change. In practice these add integral action to inputs or outputs, and tend to reduce phase margin. When augmenting the plant with integrators, we should take the opportunity to add compensating zeros that can be used to shape cross-over in order to improve transient performance and robustness. Using weights $\mathcal{W}_1$ and $\mathcal{W}_2$ to shape the frequency response does exactly this. It is an effective way to achieve the same steady-state effects while also addressing the medium-frequency performance requirements explicitly.

D. Robustifying Compensator and Closed-Loop Properties

Once $\mathcal{W}_1$ and $\mathcal{W}_2$ are designed and discretized, then the robustifying compensator $\mathcal{K}_s$ is computed for the shaped plant $\mathcal{P}_s$. This involves solving two decoupled Riccati equations, for which solutions exist under the mild conditions that the plant $\mathcal{P}$ is stabilizable and detectable, and that the weights and plant do not possess any common pole-zero cancellations. Referring to Fig. 3, the algorithm, which is

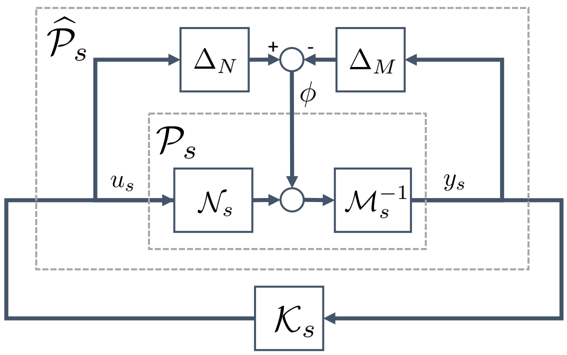

Fig. 3. The shaped plant $\mathcal{P}_s = \mathcal{M}_s^{-1}\mathcal{N}_s$, perturbed by normalized coprime factor uncertainty blocks Δ_M and Δ_N, in feedback with $\mathcal{K}_s$.

provided in the Appendix, computes the compensator $\mathcal{K}_s$ that robustly stabilizes the family of perturbed shaped plants

$$\widehat{\mathcal{P}} = \{(\mathcal{M}+\Delta_M)^{-1}(\mathcal{N}+\Delta_N) : \|\Delta_N \quad \Delta_M\|_\infty < \epsilon\} \quad (7)$$

for maximum $\epsilon > 0$, which is the stability margin, where $\mathcal{P}_s = \mathcal{M}_s^{-1}\mathcal{N}_s$ is a normalized left coprime factorization of the shaped plant $\mathcal{P}_s$ [4], [13], [14].

E. $\mathcal{H}_\infty$ Controller Realizations

The $\mathcal{H}_\infty$ loop-shaping compensator $\mathcal{K}_s$ has an observer-based state feedback structure [5], [14]. The controller can be realized in different forms depending on how the reference r enters the feedback loop. For the $\mathcal{W}_1$-integrator design in Section II-A, a common realization is

$$\begin{aligned} \widehat{x}_s(k+1) &= A_s\widehat{x}_s(k) + B_s u_s(k) \\ &\quad + H_s\left(\widehat{y}_s(k) - y_s(k)\right) && \text{(8a)} \\ \widehat{y}_s(k) &= C_s\widehat{x}_s(k) && \text{(8b)} \\ u_s(k) &= K_s\widehat{x}_s(k) + K_0 r(k), && \text{(8c)} \end{aligned}$$

where the shaped plant $\mathcal{P}_s$ is realized as

$$\begin{aligned} x_s(k+1) &= A_s x_s(k) + B_s u_s(k) + B_{qs} q(k) && \text{(9a)} \\ y_s(k) &= C_s x_s(k)\,, && \text{(9b)} \end{aligned}$$

so that x_s includes x from (4a) and states from $\mathcal{W}_1$ and $\mathcal{W}_2$, and B, B_q and C in (4a)-(4b) are suitably augmented to form B_s, B_{qs} and C_s, respectively. In (8c), the $m \times m$ matrix

$$K_0 = K_s(I - A_s - B_s K_s - H_s C_s)^{-1} H_{s1} W_{22}$$

is chosen so that $r = y_1$ for constant r in steady-state (Requirement A), $H_s = [H_{s1}\ H_{s2}]$, where H_{s1} corresponds to y_1 and H_{s2} corresponds to y_2, $W_2 = [W_{21}\ \ W_{22}]^T$ is the steady-state gain of $\mathcal{W}_2$, and W_{21} corresponds to y_1. Equations (8a)-(8b) are the state estimator with H_s being the robust estimator gain, and (8c) is the state feedback.

Alternatively, the $\mathcal{W}_1$-integrator design can be realized as

$$\begin{aligned} \widehat{x}_s(k+1) &= A_s\widehat{x}_s(k) + B_s u_s(k) && \text{(10a)} \\ &\quad + H_s\left(\widehat{y}_s(k) - y_s(k) + r(k)\right) \\ \widehat{y}_s(k) &= C_s\widehat{x}_s(k) && \text{(10b)} \\ u_s(k) &= K_s\widehat{x}_s(k) && \text{(10c)} \end{aligned}$$

where the reference r enters the controller at the input of the estimator in (10a). In this realization, the compensator has the property that for constant values of disturbance q 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 the controller $\mathcal{K}_s$.

For the $\mathcal{W}_2$-integrator design in Section II-B, the preferred realization is

$$\begin{aligned} \widehat{x}_s(k+1) &= A_s\widehat{x}_s(k) + B_s u_s(k) \\ &\quad + H_s\left(\widehat{y}_s(k) - y_s(k)\right) && \text{(11a)} \\ \widehat{y}_s(k) &= C_s\widehat{x}_s(k) && \text{(11b)} \\ u_s(k) &= K_s\widehat{x}_s(k), && \text{(11c)} \end{aligned}$$

where the shaped plant $\mathcal{P}_s$ is realized as

$$\begin{aligned} x_s(k+1) &= A_s x_s(k) + B_s u_s(k) \\ &\quad + B_{rs} r(k) + B_{qs} q(k) && \text{(12a)} \\ y_s(k) &= C_s x_s(k)\,. && \text{(12b)} \end{aligned}$$

Note that the reference r is subtracted from y_1 at the input to the weight $\mathcal{W}_2$, as shown in Fig. 2, leading to the B_{rs} term in (12a). This ensures that the integral action in $\mathcal{W}_2$ will drive the tracking error $y_1 - r$ to zero in the steady-state for constant values of q and r.

F. MPC Realization

Our objective is to construct an MPC that implements the loop-shaping controller $\mathcal{K}_s$ exactly when the constraints are inactive. We begin with the cost function. The matrices Q, R, and associated P in (2) are computed by solving an inverse-optimal control problem, which can be stated as follows. Given the shaped plant state-space model for any of the controller realizations from Section II-E,

$$x_s(k+1) = A_s x_s(k) + B_s u_s(k)\,, \quad (13)$$

and the stabilizing full state feedback

$$u_s(k) = K_s x_s(k)\,, \quad (14)$$

find a quadratic cost function

$$J(u_s) = \sum_{k=1}^{\infty} x_s^T(k) Q_s x_s(k) + u_s^T(k) R_s u_s(k)\,, \quad (15)$$

where $R_s > 0$ and $Q_s \geq 0$ such that (14) minimizes (15) subject to (13). In other words, given A_s, B_s and K_s, compute $Q_s \geq 0$ and $R_s > 0$ such that P_s and K_s satisfy the associated Riccati equation

$$A_s^T P_s A_s - A_s^T P_s B_s (R_s + B_s^T P_s B_s) B_s^T P_s A_s + Q_s = P_s \quad (16)$$

where

$$K_s = -(B_s^T P_s B_s + R_s)^{-1} B_s^T P_s A_s. \tag{17}$$

Solutions to this problem are published [15], [8], with one solution being: Set $R_s = I$, and compute Q_s by solving the same Riccati equations that are used to compute K_s and H_s in $\mathcal{K}_s$. Then (16) is used to compute P_s. Unfortunately, this approach can result in a numerically ill-conditioned P_s, which will make the resulting MPC optimization problem ill-conditioned, as reported by [12], for example. An alternative approach is to solve an LMI for $Q_s \geq 0$, $R_s > 0$ and $P_s > 0$, subject to (16) and (17), and numerically minimizing the condition number of P_s. We will assume for the remainder of this paper that a solution to the inverse problem can be computed using either approach.

Then the basic idea in LSMPC is to use the state estimator from $\mathcal{K}_s$ to initialize the states of the prediction model $\mathcal{P}_s$. However, there is a complication: The estimator states would be biased by the unmeasured disturbance q. If we use the shaped plant $\mathcal{P}_s$ as the prediction model, and initialize its state with the $\mathcal{H}_\infty$ estimator, then the constrained outputs v will be biased, leading to errors in constraint enforcement.

We emphasize that it is not feasible to augment the original plant state x with q and design the $\mathcal{H}_\infty$ estimator for the augmented plant in an attempt to integrate disturbance estimation into the $\mathcal{H}_\infty$ estimator, because such an augmented plant is not stabilizable and one of the algebraic Riccati equations in the $\mathcal{H}_\infty$ loop-shaping design cannot be solved. Furthermore, it is not feasible to add the term $B_{qs}\widehat{q}$, where $\widehat{q}$ is an estimate of the unmeasured q, to either (8b), (10a) or (11a), with the intention of removing the bias from the estimate $\widehat{x}_s$, because this will result in loss of tracking of r.

Instead, we construct an *augmented* prediction model that *includes* the shaped plant $\mathcal{P}_s$ to predict the performance output z, but is augmented with additional states to predict the constrained output v without bias. Toward this end, we construct a second estimator for q, which we denote the *disturbance estimator*. One way to do this is to rewrite (4) to include q as an additional state, assuming that it is constant over the prediction horizon,

$$\begin{bmatrix} x(k+1) \\ q(k+1) \end{bmatrix} = \begin{bmatrix} A & B_q \\ 0 & I \end{bmatrix} \begin{bmatrix} x(k) \\ q(k) \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} u(k) \tag{18a}$$

$$y(k) = [C \quad 0] \begin{bmatrix} x(k) \\ q(k) \end{bmatrix}. \tag{18b}$$

Assuming (18) is detectable, we construct a Luenberger Observer for (18), defining the process and measurement weights as, for example,

$$Q_q = \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} \qquad \text{and} \qquad R_q = \delta I, \tag{19}$$

where $\delta > 0$ is a tuning parameter, and solving the associated discrete-time Riccati equation for the gain G_q, resulting in the estimator

$$\begin{bmatrix} \widehat{x}(k+1) \\ \widehat{q}(k+1) \end{bmatrix} = \begin{bmatrix} A & B_q \\ 0 & I \end{bmatrix} \begin{bmatrix} \widehat{x}(k) \\ \widehat{q}(k) \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} u(k) + G_q\left(y(k) - \widehat{y}(k)\right) \tag{20a}$$

$$\widehat{y}(k) = [C \quad 0] \begin{bmatrix} \widehat{x}(k) \\ \widehat{q}(k) \end{bmatrix}. \tag{20b}$$

Of course the disturbance estimator design specifics depend on the particulars of an application, and other types of estimators are certainly possible. Note that this design uses the open-loop plant model $\mathcal{P}$, and is therefore sensitive to model uncertainty.

To form the prediction model, write the input weight $\mathcal{W}_1$ in state-space as

$$x_w(k+1) = A_w x_w(k) + B_w u_s(k) \tag{21a}$$

$$u(k) = C_w x_w(k) + D_w u_s(k). \tag{21b}$$

For the $\mathcal{W}_1$-integrator design, (21a) includes the integral states, while for the $\mathcal{W}_2$-integrator design, with $\mathcal{W}_1$ constant, (21) has only the direct transmission term D_w, so the dimension of x_w is zero. In either case we define the prediction model state as

$$\widehat{x}_p(k) = [\widehat{x}_s^T(k) \quad x_w^T(k) \quad \widehat{x}^T(k) \quad \widehat{q}^T(k)]^T, \tag{22}$$

the prediction model as

$$\widehat{x}_p(k+1) = \begin{bmatrix} A_s & 0 & 0 & 0 \\ 0 & A_w & 0 & 0 \\ 0 & BC_w & A & B_q \\ 0 & 0 & 0 & I \end{bmatrix} \widehat{x}_p(k) + \begin{bmatrix} B_s \\ B_w \\ BD_w \\ 0 \end{bmatrix} u_s(k) \tag{23a}$$

$$z(k) = \begin{bmatrix} I & 0 & 0 & 0 \end{bmatrix} \widehat{x}_p(k) \tag{23b}$$

$$v(k) = \begin{bmatrix} 0 & GC_w & F & 0 \end{bmatrix} \widehat{x}_p(k) \tag{23c}$$

and the cost function as

$$J(u_s) = x_s^T(N) P_s x_s(N) + \sum_{k=1}^{N-1} z^T(k) Q_s z(k) + u_s(k) R_s u_s(k). \tag{24}$$

At each sample time kT, the prediction model state (22) is initialized with $\widehat{x}_s(k)$ computed by the $\mathcal{H}_\infty$ estimator (10a)-(10b) or (11a)-(11b), and with $\widehat{x}(k)$ and $\widehat{q}(k)$ computed by the disturbance estimator (20), and with direct measurement of $x_w(k)$ (or in $\mathcal{W}_2$-integrator design, it is not present). Importantly, the performance variable z depends only on the estimated state $\widehat{x}_s$, because of the structure of E in (23b). The disturbance estimate $\widehat{q}$ affects only the predicted constrained output v and not the predicted performance output z. This ensures that the MPC is identical to the $\mathcal{H}_\infty$ loop-shaped controller, and inherits its robustness properties, when constraints are inactive.

Proposition 1: Consider the family of perturbed plants (7) and the model predictive control with prediction model (23), cost (24), estimator (20), and $\mathcal{H}_\infty$ estimator (10a)-(10b) or (11a)-(11b). Let $\mathbb{X}_{ia}$ be the set of states where the constraints (3) are strictly satisfied by the MPC input sequence along the entire prediction horizon. For every equilibrium $\hat{x}_p^e$ in the interior of $\mathbb{X}_{ia}$ there exists a set $\mathbb{S} \subseteq \mathbb{X}_{ia}$ with $\hat{x}_p^e$ in the interior of $\mathbb{S}$ such that $\mathbb{S}$ is invariant for the plant

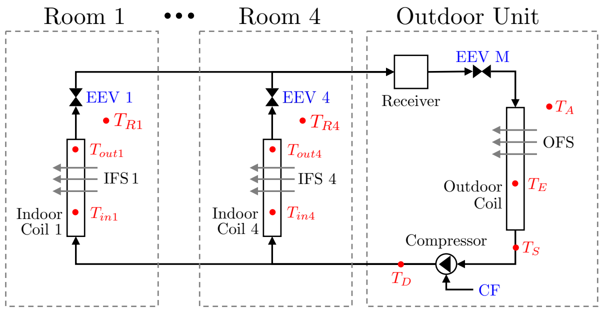

Fig. 4. Heat pump system from [16] showing the location of the temperature sensors (red) and control variables (blue). Refrigerant flows counter-clockwise.

in closed-loop with MPC, and the closed-loop is stable in $\mathbb{S}$ with the same stability margin ϵ provided by $\mathcal{H}_\infty$ loop-shaped controller corresponding to either (10) or (11).

Proof (Sketch): The proof follows directly by the design method based on inverse optimality, and hence it is only sketched. Since the MPC cost function is designed by inverse optimality, in $\mathbb{X}_{ia}$ the MPC command is equal to the command of the $\mathcal{H}_\infty$ loop-shaped controller (10) or (11). Let $\mathbb{S} \subseteq \mathbb{X}_{ia}$ be any invariant set of the closed-loop between the plant (7) and the $\mathcal{H}_\infty$ loop-shaped controller (10) or (11) with $\hat{x}_p^e$ in the interior of $\mathbb{S}$, so that starting from within $\mathbb{S}$, the closed-loop trajectory remains in $\mathbb{S}$. Since $\mathbb{S} \subseteq \mathbb{X}_{ia}$, the MPC is equal to the $\mathcal{H}_\infty$ loop-shaped controller in $\mathbb{S}$, so that $\mathbb{S}$ is invariant also for the closed-loop between the plant (7) and the MPC. Finally, the existence of $\mathbb{S}$ with the above properties follows from the closed-loop between the plant (7) and the $\mathcal{H}_\infty$ loop-shaped controller being stable, which ensures the existence of a Lyapunov function, whose sublevel sets are invariant and contain $\hat{x}_p^e$ in their interior. Hence, $\mathbb{S}$ is any such sublevel set entirely contained in $\mathbb{X}_{ia}$.

The prediction model (23) can be simplified if the problem includes constraints only on the input $u(k)$, there are no output constraints, and the $\mathcal{W}_1$-integrator design is used. Then the disturbance estimator (20) is not needed, and (22) includes only $\widehat{x}_s(k)$ and $x_w(k)$. In this case an unmeasured disturbance, if present, will be properly rejected, and input constraints properly enforced, because $x_w(k)$, which is measured, is not affected directly by $q(k)$.

III. Multi-Zone Heat Pump Case Study

Consider the heat pump shown in Fig. 4, consisting of one outdoor unit and four indoor units. The outdoor unit contains a receiver, an electronic expansion valve (EEV), denoted EEV M, an evaporating heat exchange coil, a compressor and an outdoor fan. The indoor units each contain a condensing heat exchange coil, an EEV and an indoor fan. The six controls for the system are the compressor frequency CF, the commanded settings for each EEV i, $1 \le i \le 4$, and EEV M. The fan speeds for the indoor and outdoor units are assumed constant. Each zone is subject to an unmeasured heat disturbance Q_i. The measurements are the four room temperatures T_{Ri}, the eight condenser inlet and outlet temperatures T_{ini} and T_{outi}, $1 \le i \le 4$, the evaporator temperature T_E and the compressor discharge temperature T_D.

The heat pump is modeled as a scaled discrete-time linear system $\mathcal{P}$, given in state-space form as (4a)-(4c). The model is computed by linearizing and reducing a detailed *Modelica* system model described in [17]. In addition to stability, the requirements for the closed-loop system are to

A) Track constant room temperature set-points with zero steady-state error, if possible;
B) Reject constant, unmeasured heat load disturbances, if possible;
C) Track a desired compressor discharge temperature T_D with zero steady-state error, if possible;
D) Achieve a closed-loop bandwidth ω_b for room temperature set-point tracking;
E) Enforce hard constraints on all control inputs; and
F) Enforce constraints on outputs including (1) a minimum subcooling temperature $T_{SCi} = T_{ini} - T_{outi} > T_{SCMin}$, $1 \le i \le 4$; (2) a maximum $T_D < T_{DMax}$; (3) a minimum discharge temperature super-heat $T_{DSH} = T_D - T_{ini} > T_{DSHMin}$.

We use the $\mathcal{W}_2$-integrator design with 5 regulated outputs y_1: The four room temperatures T_{Ri} and the discharge temperature error, which is the difference between T_D and a reference that is scheduled on CF, as described in [16]. PI weights are used for each of these, and constant gains are used in $\mathcal{W}_1$ and the y_2 outputs in $\mathcal{W}_2$, as shown in Fig. 2.

Frequency responses are plotted in Fig. 5. We achieve a desired cross-over frequency of approximately 0.01 rad/s, giving a time constant of about 10 min, and good phase (not shown) at cross-over, with relatively little tuning effort. The stability margin achieved is $\epsilon = 0.4$, which is excellent for this plant. In the plot of $|\mathcal{P}_s(j\omega)|$ we see that there is one singular value with a faster bandwidth, aligned strongly with the compressor input and the discharge temperature output. The weakest (lowest gain) direction is aligned in the direction of *differences* in room temperatures, which is a consequence of the heat pump architecture in which all the condensers are at the same pressure, neglecting pipe losses.

We solve the inverse optimal control problem using the method in [8]. The condition number of P_s is 10^8, which is excessive but solvable on a desktop PC. We form the prediction model, cost and constraints as described in the Section II-F. There are hard constraints on the scaled inputs, $-1.5 < u_i < 1.5$ for $1 \le i \le 6$, and soft constraints on T_{SCi} ($-10°$C), T_D ($40°$C), and T_{DSH} ($10°$C). Simulation using a linear model is performed for two different transients. For these simulations, the LSMPC horizon is set to $N = 20$, the sample period is $T = 15\,$s, making the horizon 5 minutes in length. ADMM or PQP [18] can be used to solve the optimization problem in real-time.

For the first simulation shown in Fig. 6, the temperature set-point in Room 1 is increased by $2°$C, and then the set-points in Rooms 2, 3, 4 are also increased by $2°$C, and finally a -1kW heat load step applied to all rooms. We see that LSMPC can track room temperature references with

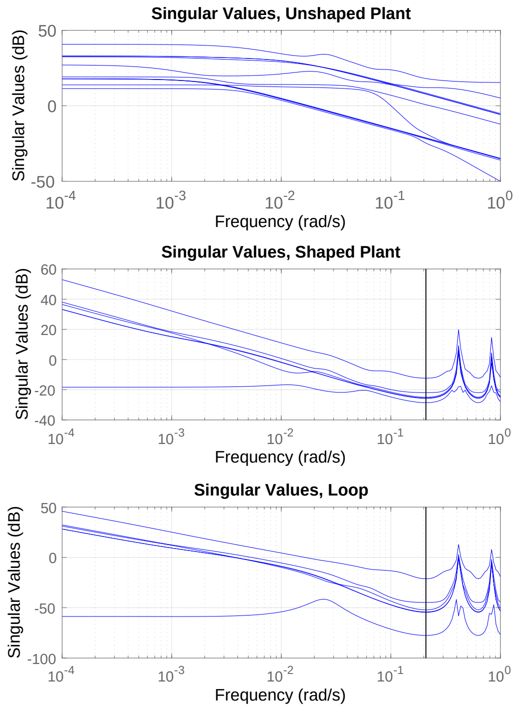

Fig. 5. Frequency response of $\mathcal{P}$ (top), $\mathcal{P}_s$ (middle), and $\mathcal{P}_s\mathcal{K}_s$ (bottom).

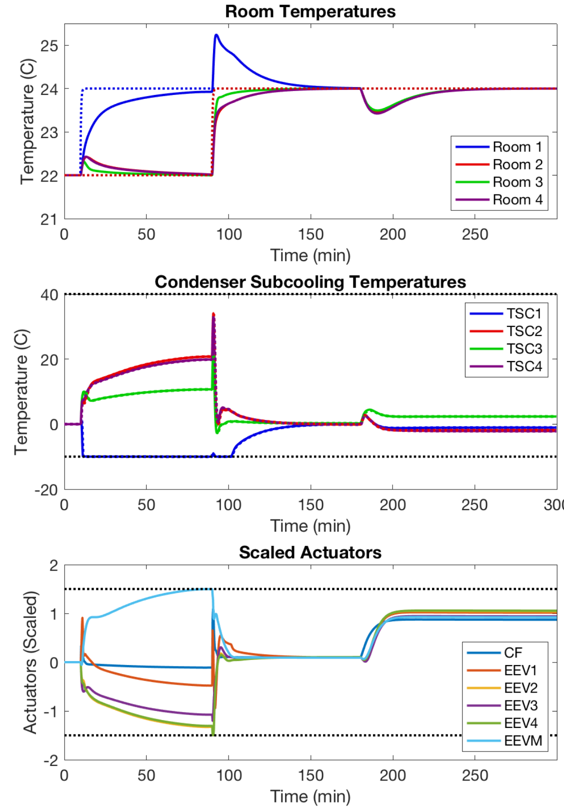

Fig. 6. Loop-shaped MPC simulation. At $t = 10\,\mathrm{min}$, the reference for Room 1 is increased $2°\mathrm{C}$, then at $t = 90\,\mathrm{min}$, references for Rooms 2, 3, 4 are increased $2°\mathrm{C}$, and finally at $t = 180\,\mathrm{min}$, a heat load step of $-1\,\mathrm{kW}$ is applied to all four rooms. Input and output constraints are active in transient, but not steady-state. References are tracked with zero steady-state error, and the heat load disturbance is rejected.

zero steady-state error when no constraints are active, reject the heat load disturbance with zero steady-state error, and enforce all constraints. Note the rise-time of the first transient is larger, because that direction (different room temperatures) is lower-gain than the second transient, where all room temperatures are the same, as expected by the frequency response analysis.

A second simulation scenario is is shown in Fig. 7. Here, the temperature set-point in Room 1 is increased by $2°\mathrm{C}$, and then the set-point in Room 2 is increased by $3°\mathrm{C}$, while Rooms 3 and 4 set-points remain at $22°\mathrm{C}$, and finally a $-1\mathrm{kW}$ heat load step applied to all rooms. This is a more aggressive case because it is in the low-gain direction of the plant. Therefore, more constraints are activated. We see that the system can track room temperatures with zero steady-state error after the first transient, but the second causes more constraints to be activated, and as a result there is a small steady-state error in the room temperatures. After the heat load transient, there is a larger steady-state error in the room temperatures, but the system is still stable, and the steady-state error is acceptable.

IV. Conclusions

An MPC formulation is presented based on $\mathcal{H}_\infty$ loop-shaping. This procedure optimizes a stability margin with respect to normalized coprime factor plant uncertainty, and provides a design methodology for the plant augmentations, cost function, and state estimator. The method can be applied to general MPC problems, but is most applicable to those problems that require state estimation. We exemplify the method for a multi-zone heat pump. The MPC for this problem has a good stability margin, enforces input and output constraints, tracks constant references and rejects constant disturbances with zero steady-state error, and provides good transient response. The tuning procedure is straightforward with a minimal number of parameters.

There are a number of extensions possible. First, alternative architectures of the $\mathcal{H}_\infty$ feedback loop, and the different ways these can be realized as an MPC should be thoroughly explored. For example, how should reference or disturbance preview be incorporated? Second, improved numerical methods for solving the inverse optimal controller by LMI should be developed. Third, it may be possible to analyze robustness for the MPC considering active constraints, since the MPC can be realized as a linear affine state feedback law for each combination of active constraints [2], and the stability margin of the feedback loop is the inverse of the $\mathcal{H}_\infty$ gain from ϕ to $[u_s^T \ y_s^T]^T$ to in Fig. 3 for the linear system [14].

Acknowledgment

We gratefully acknowledge Fumiaki Baba, Takaya Yamamoto, Masashi Fujitsuka, Yuki Mori, and Takahiro Nakai of the Advanced Technology Center, Mitsubishi Electric

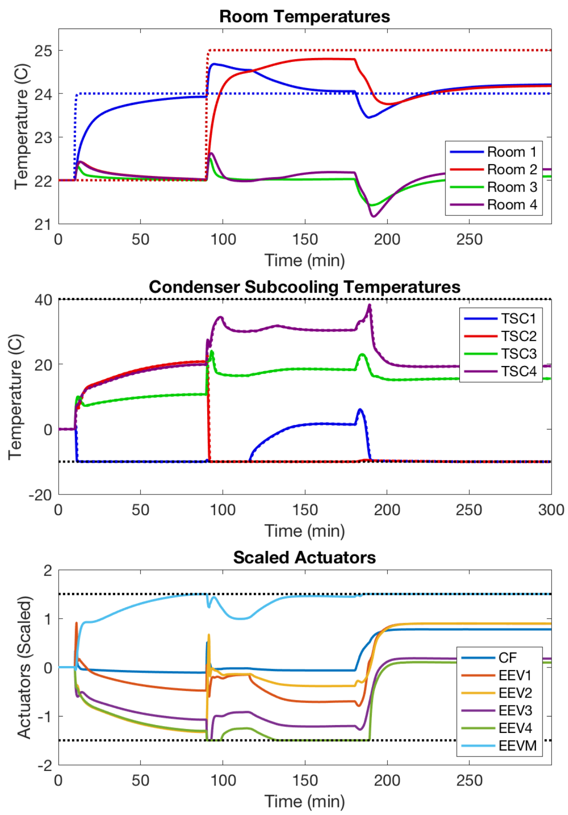

Fig. 7. A more aggressive scenario. At $t = 10\,\mathrm{min}$, the reference for Room 1 is increased $2°\mathrm{C}$, then at $t = 90\,\mathrm{min}$, reference for Room 2 is increased $3°\mathrm{C}$, while references for Rooms 3 and 4 remain at $22°\mathrm{C}$, and finally at $t = 180\,\mathrm{min}$, a heat load step of $-1\,\mathrm{kW}$ is applied to all four rooms. Input and output constraints are active in transient and steady-state, and the rooms are unable to track with zero steady-state error, but constraints are enforced, tracking errors are relatively small and the closed-loop remains stable.

Corporation, Japan, for their strong collaboration and unwavering support of this work.

References

[1] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation and Design, 2nd Edition*. Nob Hill Publishing, LLC, 2017.

[2] F. Borrelli, A. Bemporad, and M. Morari, *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.

[3] J. Doyle, "Guaranteed margins for LQG regulators," *IEEE Transactions on Automatic Control*, vol. 23, no. 4, pp. 756–757, Aug 1978.

[4] 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.

[5] J. Sefton and K. Glover, "Pole/zero cancellations in the general H_∞ problem with reference to a two block design," *Systems & Control Letters*, vol. 14, no. 4, pp. 295–306, April 1990.

[6] J. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optimization Methods and Software*, vol. 11–12, pp. 625–653, 1999, version 1.05 available from `http://fewcal.kub.nl/sturm`.

[7] J. Löfberg, "YALMIP : A toolbox for modeling and optimization in MATLAB," in *In Proceedings of the IEEE CACSD Conference*, Taipei, Taiwan, 2004.

[8] C. Rowe and J. Maciejowski, "Tuning MPC using H_∞ loop shaping," in *Procedings of the American Control Conference*, 2000, pp. 1332–1336.

[9] J. Maciejowski, "Reverse-engineering existing controllers for MPC design," *IFAC Proceedings*, vol. 40, no. 20, pp. 436–441, 2007.

[10] S. Di Cairano and A. Bemporad, "Model predictive control tuning by controller matching," *IEEE Transactions on Automatic Control*, vol. 55, no. 1, pp. 185–190, 2010.

[11] B. P. Molinari, "The stable regulator problem and its inverse," *IEEE Transactions on Automatic Control*, vol. AC-18, no. 5, pp. 454–459, October 1973.

[12] M. C. Priess, R. Conway, J. Choi, J. John M. Popovich, and C. Radcliffe, "Solutions to the inverse LQR problem with application to biological systems analysis," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 2, pp. 770–777, March 2015.

[13] G. Vinnicombe, *Uncertainty and Feedback: H_∞ Loop-Shaping and the ν-Gap Metric*. Imperial College Press, 2001.

[14] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*. Wiley, 2005.

[15] P. A. Iglesias, "The strictly proper discrete-time controller for the normalized left-coprime factorization robust stabilization problem," *IEEE Transactions on Automatic Control*, vol. 45, no. 3, pp. 516–520, March 2000.

[16] S. A. Bortoff, D. J. Burns, C. R. Laughman, H. Qiao, C. Danielson, A. Goldsmith, and S. Di Cairano, "Power optimizing control of multi-zone heat pumps," in *IEEE Conference on Control Technology and Applications*, Aug. 2018, pp. 826–833.

[17] 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.

[18] S. Di Cairano, M. Brand, and S. A. Bortoff, "Projection-free parallel quadratic programming for linear model predictive control," *International Journal of Control*, vol. 86, no. 8, pp. 1367–1385, Aug. 2013.

Appendix

In this Appendix the equations for calculation of the $\mathcal{H}_\infty$ loop-shaping controller and inverse optimal controller from [8] are included for completeness. For the strictly proper shaped plant

$$x_s(k+1) = A_s x_s(k) + B_s u_s(k) \tag{25}$$

$$y_s(k) = C_s x_s(k) \tag{26}$$

compute solutions to the pair of Riccati equations

$$A_s^T X_s B\left(B_s^T X_s B_s + I\right) B_s^T X_s A_s - A_s^T X_s A_s + X_s = C_s^T C_s \tag{27}$$

$$A_s Z_s C_s^T\left(C_s Z_s C_s^T + I\right) C_s Z_s A_s^T - A_s Z_s A_s^T + Z_s = B_s B_s^T. \tag{28}$$

Stabilizing solutions to (27) and (28) exist if (A_s, B_s) is stabilizable and (A_s, C_s) is detectable, which is ensured by conventional loop-shaping techniques. Next compute

$$\gamma_{min} = \sqrt{1 + \lambda_{max}(X_s Z_s)}, \tag{29}$$

where λ_{max} is the maximum eigenvalue. Then, given a value $\gamma_{rel} > 1$, which is the relative sub-optimality (typically $\gamma_{rel} = 1.1$), compute $\gamma = \gamma_{rel} \cdot \gamma_{min}$. Then

$$W_s = (\gamma^2 - 1)I - Z_s X_s \tag{30}$$

$$K_s = -\gamma^2 B_s^T X_s W_s^{-1}\left(I + \gamma^2 B_s B_s^T X_s W_s^{-1}\right)^{-1} A_s \tag{31}$$

$$H_s = -A_s Z_s C_s^T\left(I + C_s Z_s C_s^T\right)^{-1} \tag{32}$$

$$Q_s = \gamma^2 X_s W_s^{-1} - \gamma A_s^T\left(I + \gamma^2 X_s W_s^{-1} B_s B_s^T\right)^{-1} X_s W_s^{-1} A_s \tag{33}$$

$$R_s = I. \tag{34}$$

and $\epsilon = 1/\gamma$.