Approaching simulation-based controller design: heat exchanger case study

Matias Waller[a,b,*]   Leonardo Espinosa-Leal[b]

a Åland University of Applied Sciences, Mariehamn, Åland; b Arcada University of Applied Sciences, Helsinki, Finland
[*]matias.waller@arcada.fi

Abstract

Supported by an identification experiment using random-phase multisines, uncertain parameters in a grey-box model for a multiple-input multiple-output laboratory-scale heat exchanger are fitted to experimental data. By defining desired trajectories for the controlled system concerning setpoint changes, simulations and a cost function taking control signal activity into account, we determine both a linear and a nonlinear PI-controller. The resulting control systems are evaluated through practical experiments and analysis with encouraging results. The approach to modelling and controller design raises questions about what is needed from an educational point of view, e.g., what skills are needed for simulation-based control design and analysis?

1 Introduction

The development of increasingly accurate models and simulations raises questions about whether new approaches to the design of control systems should be considered. As an example, models often popularly referred to as Digital Twins (DTs) (Espinosa-Leal et al., 2020) that combine sensor information, data-based black-box models with physical models to develop a faithful virtual replica of a given system are used onboard in the electronic control unit for improving the final tuning phase of the controller for the physical vehicle in an automobile application (Dettù et al., 2023). Inspired by similar approaches and the approach of (nominal) model predictive control (mpc) in general (Hewing et al., 2020), we explore possibilities to use simulations directly for controller design. Our case study considers a laboratory-scale heater and a heat exchanger. Heat exchangers are vital components in, among other applications, thermal power stations and heat pumps. Improved control of heat exchangers based on the creative use of increasingly accurate models could thus contribute to the much needed increased energy efficiency (IPCC, 2023).

Although it would be appealing to illustrate the approach on a deterministic, linear model that could be fully explored analytically, the primary motivation for using simulations directly is, naturally, for processes with features not easily captured by linearized models. Therefore, the approach is illustrated in a nonlinear, multivariable process. In this initial approach, we, for manageability, focus on the control of one vital quality variable. Combining first-principles with ex-periments, a system of coupled nonlinear differential equations forms the basic model of the heating process considered. Some uncertain parameters of the model are fitted based on an identification experiment. A sequence of setpoint changes and a control performance criterion are then defined. Different control strategies are explored and optimized based on simulations. The strategies are then implemented in practice, results are evaluated, and a preliminary analysis of the control systems is presented.

Reflecting on the approach, what does it imply regarding modelling and model analysis required for teaching control engineering on a general level? If control design is based on simulations only, does this mean that the skills needed for analyzing (and linearizing) differential or difference equations are less critical? How can traditional requirements on stability analysis, robustness, control performance and control signal activity be explored based on simulations? Connected to these questions is the opaque nature of general nonlinear black-box models and physical component models with block diagrams that often make them less suitable for traditional, linear approaches. The gaining popularity of these modelling alternatives also motivates the approach presented in this paper.

2 Laboratory-scale heat exchanger

A schematic view of the process is illustrated in the left panel of Fig. 1. The process has three control signals, u_c controlling the pump on the cold side, u_h controlling the pump on the hot side, and u_p controlling the power to the heater. Given the external disturbances

T_{ci} and T_{surr}, i.e., the temperature of the flow on the cold side into the heat exchanger and the surrounding temperature, a model should be able to determine the power to the heater P, the flows on the cold and on the hot side, $\dot{V}_c$ and $\dot{V}_h$, the temperatures entering and exiting on the hot side, T_{hi} and T_{ho}, and the temperature exiting on the cold side, T_{co}.


Figure 1. Schematic view of the heating process. Blue denotes the cold side of the heat exchanger, and red denotes the hot side and the heater.

The main objective is to control $\dot{V}_c$ as well as T_{co}. The temperatures on the hot side are less interesting, but naturally sufficient heat is necessary to enable control of T_{co}.

2.1 A grey-box model

In the approach to the simulation-based design of controllers, we consider a simplified first-principles model for the heating process. A net energy balance can be written

$$\frac{dE_{tot}}{dt} = P - \dot{V}_c \rho c_p (T_{co} - T_{ci}) - \dot{Q}_{loss} \quad (1)$$

where dE_{tot}/dt is the change of net stored energy in the system and $\dot{Q}_{loss}$ are the heat losses to the environment. The term E_{tot} cannot easily be used to determine the temperatures T_{hi}, T_{ho} and T_{co}. For this reason, the net balance is split into three equations. A balance over the heater yields

$$\frac{dE_p}{dt} = P - \dot{V}_h \rho c_p (T_{hi} - T_{ho}) - \dot{Q}_{loss} \quad (2)$$

where all heat losses from the system are assigned to the heater. The stored energy is assumed to be characterized by the temperature T_{hi} and given by $E_p = C_p(T_{hi} - T_{ref})$ where C_p is the heat capacity for the

heater and T_{ref} is a reference temperature. If C_p is assumed constant, $dE_p/dt = C_p dT_{hi}/dt$. A similar balance over the hot side of the heat exchanger yields

$$\frac{dE_h}{dt} = \dot{V}_h \rho c_p (T_{hi} - T_{ho}) - \dot{Q}_{he} \quad (3)$$

where $\dot{Q}_{he}$ is the power transferred in the heat exchanger and dE_h/dt is characterized by T_{ho}. If C_h is the heat capacity for the hot side of the heat exchanger and assumed constant, $dE_h/dt = C_h dT_{ho}/dt$. The corresponding balance over the cold side of the heat exchanger is given by

$$\frac{dE_c}{dt} = \dot{Q}_{he} - \dot{V}_c \rho c_p (T_{co} - T_{ci}) \quad (4)$$

where dE_c/dt is characterized by the temperature T_{co}. If C_c is the heat capacity for the cold side of the heat exchanger and assumed constant, $dE_c/dt = C_c dT_{co}/dt$.

In addition to these energy balances, equations for $\dot{Q}_{loss}$, $\dot{Q}_{he}$ and equations for dependencies between u_c and $\dot{V}_c$, u_h and $\dot{V}_h$ and u_p and P are needed.

For a heat exchanger, it is common to use $\dot{Q}_{he} = \alpha A \Delta T_{lm}$, where α is the heat transfer coefficient, A is the exchange area and ΔT_{lm} is the logarithmic mean temperature difference. In this case, the temperature differences are $T_{hi} - T_{co}$ and $T_{ho} - T_{ci}$. Although ΔT_{lm} is motivated by steady-state, we nonetheless use it for our dynamic model. The heat losses are assumed to be proportional to $T_{hi} - T_{surr}$, i.e., $\dot{Q}_{loss} = k(T_{hi} - T_{surr})$. For the pumps and the heater, first-order models were fitted to step experiments to give the equations,

$$\frac{d\dot{V}_c}{dt} = \frac{1}{T_c}(K_c \max(u_c - 19, 0) - \dot{V}_c)$$

$$\frac{d\dot{V}_h}{dt} = \frac{1}{T_h}(K_h \max(u_h - 19, 0) - \dot{V}_h) \quad (5)$$

$$\frac{dP}{dt} = \frac{1}{T_p}(K_p u_p - P)$$

with $T_c = 1.5$ s, $K_c = 5$ (ml/(min)/%), $T_h = 1$ s, $K_h = 10$ (ml/(min)/%), $T_p = 1$ s and $K_p = 0.016$ kW/%.

The models include several uncertain parameters in addition to the pump and power characteristics given above, mainly α, k, C_c, C_h, and C_p. Constant estimates for these will, for simplicity, be considered. In addition, delays are determined by visual inspection of step changes and possible flow-dependent transport delays as well as distribution of temperatures and flows, ageing and other time-variant characteristics are neglected. The identification experiment and how the parameters α, k, C_c, C_h and C_p are fitted is described next.

2.2 Identification experiment

Although α and k could be fitted to steady-state data, estimating heat capacities requires experiments with

nonzero temperature derivatives. In future work, we also aim to explore various data-based models for simulation-based controller design. For this reason, a 6000 seconds identification experiment using random-phase multisines (Pintelon & Schoukens, 2012) for all three control signals, u_c, u_h and u_p, was conducted. The main experiment contains two independent sets of three 1000 seconds long periods each. A short segment of the experiment illustrating T_{co}, corresponding simulations $T_{co,s}$ and the three control signals are illustrated in Fig. 2.


Figure 2. Upper panel: Simulations of $T_{co,s}$ (blue) and measurements T_{co} (red) for a short segment of the identification experiment. Lower panel: The three control signals for the corresponding segment of the experiment.

As Fig. 2 shows, significant variations in, especially, u_c and u_h have been implemented. Consequently, the experiment should provide valuable data over a large range. To fit the uncertain parameters to the data, simulations are compared to measurements and a quadratic criterion defined by Eq. 6 is introduced.

$$V_f = \frac{1}{m} W^T \operatorname{diag}(X^T X) W \qquad (6)$$

In Eq. 6, m is the number of observations, W is a weight vector that emphasizes the chosen variables, and X is a matrix with the simulation errors of interest as columns. The k-th row in X has the four columns $T_{co,s}(k) - T_{co}(k)$, $T_{hi,s}(k) - T_{hi}(k)$, $T_{ho,s}(k) - T_{ho}(k)$ and $\dot{V}_{c,s}(k) - \dot{V}_c(k)$ and $1 \leq k \leq m$. Controlling T_{co} is the main objective in our application and motivates our choice of $W^T = \begin{pmatrix} \sqrt{10} & 1 & 1 & 1 \end{pmatrix}$.

In addition to the main experiment, experimental step changes in the control signals were performed to estimate delays. Based on visual inspection, delays from the control signals were estimated as 2, 2, and 6 seconds respectively. In other words, $u_c(t-2)$, $u_h(t-2)$ and $u_p(t-6)$ replaces u_c, u_h and u_p in the equations for the grey-box model. It can be noted that these estimates are somewhat arbitrary, e.g., is the delay from u_h to T_{co}, to T_{hi} or to T_{ho}? Maintaining the physical interpretation of the grey-box model, we choose delays to $\dot{V}_c$, $\dot{V}_h$ and P, respectively. Even so, variations be-

tween different step changes can be discerned. Moreover, for data-based models, other delays can be motivated.

Using constrained optimization, the estimates $\alpha = 0.30\,\text{kW}/(\degree\text{Cm}^2)$, $k = 0.0033\,\text{kW}/\degree\text{C}$, $C_c = 0.17\,\text{kJ}/\degree\text{C}$, $C_h = 0.13$ kJ/$\degree$C and $C_p = 5.9$ kJ/$\degree$C are obtained. Simulations of T_{co} denoted $T_{co,s}$ for the fitted model are illustrated along with measurements in Fig. 2.

3 Designing controllers

Traditionally, controller design is based on linear models of the dynamical systems. Rules-of-thumb approaches are based on simple models. With more detailed (linear) models, the design typically addresses one of, e.g., a desired stability margin, control performance as quantified by quadratic costs in control error and control signal activity, disturbance rejection, robustness by guaranteeing stable control under uncertainties, etc. Typically, these approaches require tools for differential and difference equations, linearizing equations, state-space descriptions and linear algebra, block diagrams, frequency analysis and Bode-diagrams, optimal control, etc. Accordingly, acquiring such skills forms a major focus of control engineering courses.

With increasingly accurate models, it appears that approaches based on linearized models do not take full advantage of available insights. Furthermore, maybe nonlinear approaches to control based on local linearity, such as gain scheduling, primarily are the result of adapting the design of controllers to traditional approaches?

As an alternative, mpc is not based on designing a static control law. Instead, control signal sequences are determined by optimizing simulations of a model to follow a desired trajectory over a predictive horizon. With new measurements, optimal control signal sequences are updated based on the available state. This approach has many attractive features, e.g., nonlinear models and constraints can easily be included. A disadvantage is that it may be difficult to determine the required computational complexity a priori, e.g., hard nonlinearities, bifurcations, etc., can render optimization unfeasible. Thus, simplifications (linearizing) can be needed to guarantee necessary computational efficiency.

In this paper, we instead combine the approach of designing a static control law with that of using simulations. With a static control law, the need for computational power in the real-time implementation is negligible and this is a key motivation behind our approach and a clear advantage compared to nominal mpc. The parameters in the controller are determined by optimization, i.e., similarly to mpc we formulate and minimize a criterion that quantifies differences in simulations of process values from desired setpoint trajecto-

ries taking control signal activity into consideration. The general criterion is similar to Eq. (6) and is given by

$$V_c = \frac{1}{m}(Q^T \operatorname{diag}(X_c^T X_c)Q + R^T \operatorname{diag}(\Delta u^T \Delta u)R \quad (7)$$

where X_c is given by differences between setpoint (r) and simulated process value (y) and its k-th row is given by

$$X_c(k) = \begin{pmatrix} r_1(k) - y_1(k) & r_2(k) - y_2(k) & \cdots \end{pmatrix} \quad (8)$$

In our quadratic criterion, control signal activity is quantified by the change in control signal, $\Delta u = u(k+1) - u(k)$, but u can naturally be used directly if more suitable for the application at hand.

The idea of simulation-based controller design is not new, e.g., neural network controllers can often be seen as a version of this approach. But the idea and criterion has to the best of our knowledge not been as explicitly discussed as we try in this paper. Furthermore, we also discuss how to incorporate feedback systems analysis within our framework.

For our specific case study, the focus of the present study is to explore possibilities for controlling T_{co} using u_h. In Eq. (7), $X_c = T_{co,sp} - T_{co,s}$ and $\Delta u = \Delta u_h$ accordingly, and with the heuristic choices of $Q = 20\sqrt{2}$ and $R = 1$ based on evaluations of simulations. For the other control variables, the flow on the cold side, $\dot{V}_c$, with the pump on the cold side, u_c, and the temperature entering the heat exchanger on the hot side, T_{hi}, with u_p, we use fixed PI-controllers.

For the PI-controllers, we use the velocity equation

$$u_i(k) = u_i(k-1) + K_i\left(\left(1 + \frac{T_s}{T_{I,i}}\right)e_i(k) - e_i(k-1)\right) \quad (9)$$

with control signals constrained between 0 and 100%, $u_i(k) = \min(\max(u_i(k), 0), 100)$. In Eq. (9), K is the (proportional) gain, T_I the integration time, T_s the sampling period ($T_s = 1$ s in all simulations and experiments) and $e_i(k) = r_i(k) - y_i(k)$ is the control error, i.e., the difference between setpoint and process value. The subscript i is either c for the pump on the cold side, h for the pump on the hot side or p for the command to the power in the heater. The PI-controller for u_c has $K_c = 0.01$ %/(ml/min) and $T_{I,c} = 0.5$ s and u_p has $K_p = 20$ %/°C and $T_{I,p} = 100$ s[1].

To explore realistic challenges, the setpoint for $\dot{V}_c$ changes from 150 ml/min to 220 ml/min at $t = 1000$ s as a ramp stretching over 5 seconds and then to 100 ml/min at $t = 2000$ s as a ramp stretching over 10 seconds. These setpoint changes can be seen in Fig. 3. The setpoint for T_{co} changes from 35 °C to 40 °C at $t = 500$ s, back to 35 °C at $t = 1500$ s and to 40 °C at

---

[1] These choices are at least not intentionally biased for our study as they are arbitrarily chosen from submitted student assignments.

$t = 2500$ s. All changes are ramp-shaped and stretch over 10 seconds. These setpoint changes can be seen in Figs. 4 and 6. The setpoint for T_{hi} is kept constant at 55 °C.

3.1 PI control

Minimizing Eq. (7) for the simulated response to the setpoint changes regarding K_h and $T_{I,h}$ gives the results $K_h = 4.7$ %/°C and $T_{I,h} = 23$ s. Corresponding simulations are illustrated in the left panels of Figs. 3–5. For reference and comparison, the value for our criterion of Eq. (7) is $V_{c,e} = 510$. In Figs. 3–5, experimental results using the same setpoint changes are illustrated in the right panels. In addition to the presence of measurement noise, some differences can be noted. In the simulations, T_{co} does not reach the setpoint after the change in $\dot{V}_c$ at $t = 1500$ and u_h saturates at 100%. In the experiment, this disturbance in T_{co} is significantly smaller and compensated for by the controller with u_h just under 70%. This discrepancy between model and experiment could be explained by a higher efficiency of the heat exchanger at higher flows, i.e., α could better be described as a function of flows and, possibly, temperatures. For the control, however, it appears as if the well-known strength of integral action in the controller is robust against such low-frequency modelling errors. Perhaps more interesting, the disturbance in T_{co} at $t = 2000$ due to the change in $\dot{V}_c$ causes slowly converging oscillations, suggesting a bifurcation in the simulated closed-loop system that could warrant further study. As an advantage in practice, we note that control of T_{co} is similar but better in experiments than in simulations.

For the control of T_{hi} much larger high-frequency variation in both T_{hi} and u_p can be noted in the experiment compared to simulations. Although the control on a general level is similar in simulation and experiment and works well since high-frequency variations in u_p are not a matter of concern, this discrepancy will be briefly discussed in Section 4.

For the experiment, the value $V_{c,e} = 140$ is obtained. This significantly lower value despite measurement noise compared to the simulation can partly be explained by the simulated T_{co} not reaching the setpoint for $1500 \leq t \leq 2000$ as noted before.

3.2 Nonlinear PI control

In addition to enabling the use of detailed and, possibly, opaque models, one of the points behind using simulations for controller design is the possibility for exploring alternative, nonlinear controllers. As the left panels in Figs. 4–5 clearly reveal, the model exhibits challenging nonlinear characteristics. An intuitive solution could therefore be to consider a nonlinear PI-controller. Instead of gain-scheduling with

Figure 3. Upper panels: Setpoint changes in of $\dot{V}_c$ (blue) and simulated control (red) in the left panel, and the corresponding experiment to the right. Lower panels: Simulated control signal u_c (left)


Figure 4. Upper panels: Setpoint changes in of T_{co} (blue) and simulated control (red) in the left panel, and the corresponding experiment to the right. Lower panels: Simulated control signal u_h (left) and corresponding experiment (right).

different linear PI-controllers depending on a parameter (typically u, y or r), we explore linear dependencies in the gain and integration time. Specifically, in place of K_h and, $T_{I,h}$ we use $K_h + au_h(k-1)$ and $T_{I,h} + bu_h(k-1)$.

Minimizing Eq. (7) for the simulated response to the setpoint changes regarding K_h, a, $T_{I,h}$ and b gives the results $K_h = 0.1$ %/°C, $a = 46$ 1/°C, $T_{I,h} = 0.18$ s and $b = -0.46$ s/%. Compared to the linear PI-controller of Eq. (9) this corresponds to a range of values for K_h between 3.5 %/°C and 18 %/°C and for $T_{I,h}$ between 0.14 s and 38 s. Simulations yield the value $V_{c,e} = 240$, i.e., a significant improvement over the linear case.

The experiment with the nonlinear PI-controller yields the value $V_{c,e} = 52$. Compared to the experiment with the linear controller, the performance is numerically clearly better. Visual inspection of Figs. 6–7 reveal that setpoint tracking is faster and disturbance rejection better than in the linear case, at the cost of higher activity in the control signal u_h. In summary, the results are very encouraging and a significant improvement in control quality can be noted.
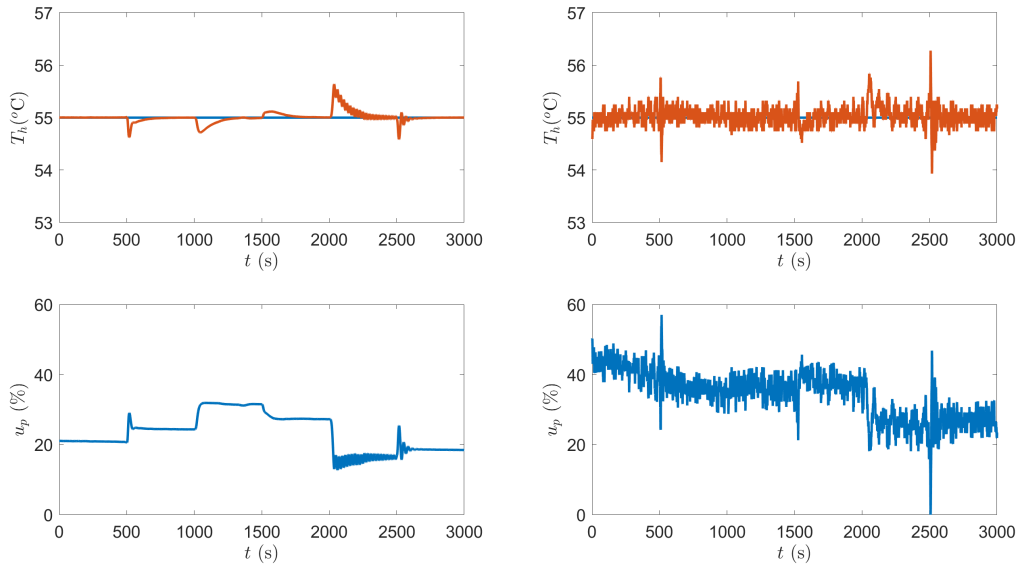
Figure 5. Upper panels: Setpoint for T_{hi} (blue) and simulated control (red) in the left panel, and the corresponding experiment to the right. Lower panels: Simulated control signal u_p (left) and corresponding experiment (right).


Figure 6. Same as in Fig. 4 but for the nonlinear PI-controller.

4   Control system analysis

It is commonly known that control design methods "focus on one or two aspects of the [control] problem, and the control-system designer then has to check that the other requirements are also satisfied" (Åström & Wittenmark, 1997). Using a block diagram and transfer function notation with s as the Laplace variable, a feedback system following the structure we use is depicted in Fig. 8.

Correspondingly,

$$Y(s) = \frac{G_r G_p}{1 + G_r G_p} R(s) + \frac{G_p}{1 + G_r G_p} W_1(s) + \frac{1}{1 + G_r G_p} W_2(s) \quad (10)$$

and

$$U(s) = \frac{G_r}{1 + G_r G_p} R(s) - \frac{G_r G_p}{1 + G_r G_p} W_1(s) - \frac{G_r}{1 + G_r G_p} W_2(s) \quad (11)$$
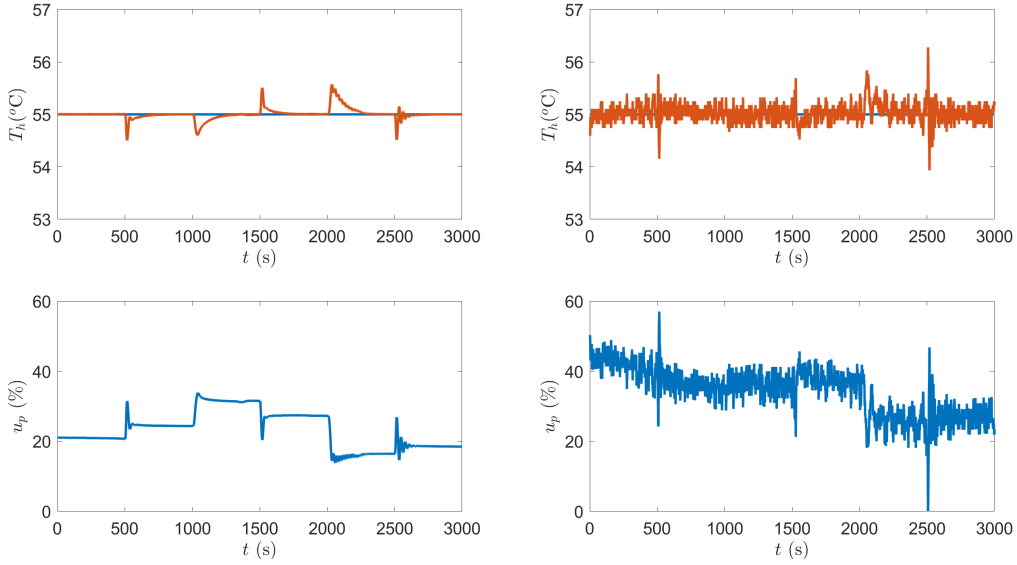
**Figure 7.** Same as in Fig. 5 but for the nonlinear PI-controller.

Figure 8. Block diagram for a process $G_p(s)$ controlled by (negative) feedback. The diagram includes setpoint $R(s)$, control signal $U(s)$, process disturbance $W_1(s)$, measurement noise $W_2(s)$, process value $Y(s)$ and controller $G_r(s)$.

As these equations indicate, the control problem can, in turn, often be analyzed (in the frequency domain) by addressing the "Gang Of Four" (Åström & Murray, 2021), i.e, the four transfer functions

$$\frac{G_r G_p}{1 + G_r G_p} \quad (12)$$

known as the complementary sensitivity function,

$$\frac{G_p}{1 + G_r G_p} \quad (13)$$

the load sensitivity function,

$$\frac{1}{1 + G_r G_p} \quad (14)$$

the sensitivity function and

$$\frac{G_r}{1 + G_r G_p} \quad (15)$$

the noise sensitivity function. For linear systems, it is illustrative to analyze the closed-loop system with plots of (the gains of) these transfer functions as a function of frequency. In principle, simulations using sinusoidal functions of different frequencies for, e.g., $w_2(t)$ and recording the corresponding amplitudes of $y(t)$ and $u(t)$ could provide numerical estimates of these gains. However, nonlinear systems can exhibit, e.g., frequency spreading, i.e., a single frequency in $w_2(t)$ can result in several frequencies in $y(t)$ and $u(t)$. This phenomenon can further be amplitude-dependent and, for multivariable systems, the principle of super-position is not necessarily applicable. In addition, the behavior can depend on the region of operation.

In summary, these characteristics render a full exploration of our case study cumbersome at the very least. We are still grasping at how the abundance of combinations of different variables could be illustrated. As an initial exploration, and motivated by the challenges the experiments indicate, we separately explore how sinusoidal measurement noise in T_{co} and T_{hi} for different frequencies affect T_{co}, u_h, T_{hi} and u_p, respectively. Constant setpoints are used, 37.5°C for T_{co} and

160 ml/min for $\dot{V}_c$. The amplitude of the measurement noise was 1°C and possible frequency spreading was neglected, i.e., only amplitudes of corresponding oscillations were recorded. The corresponding results are illustrated in Figs. 9–10.
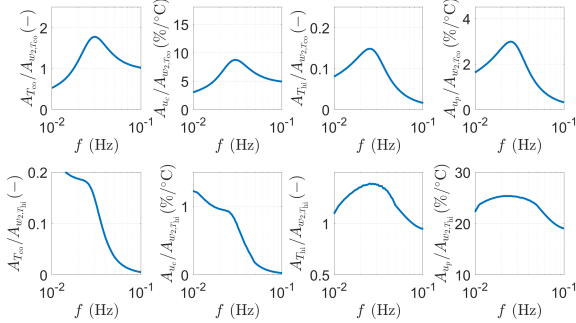

Figure 9. Amplitude of variations as a function of frequency resulting from a sinusoidal measurement disturbance on T_{co} (upper panels) with amplitude 1°. Left to right, amplitudes of T_{co}, u_c, T_{hi} and u_h. Lower panels: Same as upper panels but for T_{hi} with amplitude 1°.
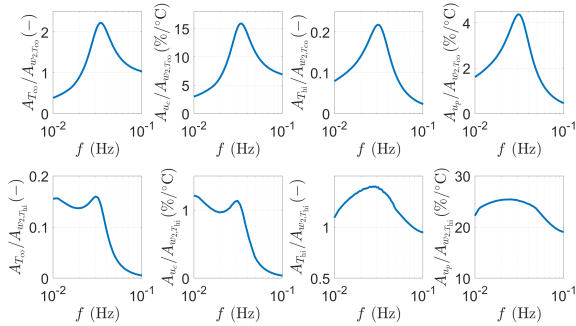

Figure 10. Same as in Fig. 9 but for the nonlinear PI-controller.

As the figures reveal, similar "sensitivity functions" are obtained for the linear and nonlinear PI-controllers, indicating a (limited) insensitivity to process variations. Based on this limited analysis, the main source of possible concern is in the "noise sensitivity functions": variations in T_{hi} can be amplified to the control signal u_p by a factor of 25 and to u_h by a factor of 3 using the linear PI-controller and a factor of 4 using the nonlinear PI. Variations in T_{co}, in turn, can be amplified to the control signal u_h by a factor of 9 (linear PI) and by 16 (nonlinear PI).
This analysis can explain the control signal activity observed in u_h (lower right panels of both Fig. 5 and Fig. 7) as well as the difference between control signal activity observed in u_p (lower right panels of Fig. 4 and Fig. 6). Most likely, the high gain of 25 is due to the heuristically designed PI-controller for controlling T_{hi} using u_p and it can be noted that the gain

$K_p = 20\,\%/°\text{C}$ was used. Still, a comparative study of different approaches for controller design, possibly including simulation-based loop-shaping, could help to shed light on possible weaknesses in the simulation-based approach. Moreover, an interesting possibility could be to include, e.g., the maximum gains of noise sensitivity functions in the criterion minimized for designing controllers, Eq. (7).

5 Conclusions and future work

In this paper, we explored simulation-based controller design. Compared to approaches based on linearized models, this enabled us to explore and tune alternative nonlinear controllers. Moreover, the approach does not rely on models that are translucent and easily linearized. For our challenging case study on a multivariable, nonlinear heating system, the simulations were promising and experimental results above expectations. Further experiments, study, analysis, and comparisons to other promising frameworks such as reinforcement learning or mpc will be pursued in order to explore the general applicability of the presented approach.

References

Åström, K. J., & Murray, R. M. (2021). *Feedback systems: An introduction for scientists and engineers* (2nd ed.). Princeton University Press.

Åström, K. J., & Wittenmark, B. (1997). *Computer controlled systems: Theory and design* (3rd ed.). New Jersey: Prentice-Hall.

Dettù, F., Formentin, S., & Savaresi, S. M. (2023). The twin-in-the-loop approach for vehicle dynamics control. *arXiv preprint arXiv:2209.02263*.

Espinosa-Leal, L., Chapman, A., & Westerlund, M. (2020). Autonomous industrial management via reinforcement learning. *Journal of intelligent & Fuzzy systems*, *39*(6), 8427–8439.

Hewing, L., Wabersich, K. P., Menner, M., & Zeilinger, M. N. (2020). Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, *3*, 269–296.

IPCC. (2023). Summary for policymakers [Book Section]. In H. Lee & J. Romero (Eds.), *Climate change 2023: Mitigation of climate change. contribution of working groups i, ii, iii to the sixth assessment report of the intergovernmental panel on climate change.* Geneva Switzerland: IPCC, in press.

Pintelon, R., & Schoukens, J. (2012). *System identification—a frequency domain approach* (2nd ed.). New Jersey: John Wiley & Sons.