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

Robust Unfalsified Control of a Heat Pump

Bortoff, Scott A.; Tsuji, Kosei

TR2025-128 September 03, 2025

Abstract

Modern heat pumps operate over increasingly wide ranges of temperatures, heat loads, compressor and fan speeds, making robustness of feedback controllers an increasingly critical issue. This paper presents results of a model-based analysis of process gain variation, and identifies a relatively simple approach for scheduling compensator gains for a conventional heat pump feedback architecture. The approach is effective at addressing process gain variation caused by widely ranging values of compressor and fan speeds. To address potential loss of feedback control robustness caused by plant uncertainty that is unmeasured, typically due to heat pump installation and application variability, the theory of unfalsified control is studied. An unfalsified control algorithm is designed for one feedback loop that is particularly susceptible to unmeasured plant uncertainty, and shown to be functional in a nonlinear simulation. Several potential issues associated with maturing unfalsified control theory into a practical technology for application to production heat pumps are discussed.

IEEE Conference on Control Technology and Applications (CCTA) 2025

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

Robust Unfalsified Control of a Heat Pump

Scott A. Bortoff [1]and Kosei Tsuji[2]

Abstract— **Modern heat pumps operate over increasingly wide ranges of temperatures, heat loads, compressor and fan speeds, making robustness of feedback controllers an increasingly critical issue. This paper presents results of a model-based analysis of process gain variation, and identifies a relatively simple approach for scheduling compensator gains for a conventional heat pump feedback architecture. The approach is effective at addressing process gain variation caused by widely ranging values of compressor and fan speeds. To address potential loss of feedback control robustness caused by plant uncertainty that is unmeasured, typically due to heat pump installation and application variability, the theory of unfalsified control is studied. An unfalsified control algorithm is designed for one feedback loop that is particularly susceptible to unmeasured plant uncertainty, and shown to be functional in a nonlinear simulation. Several potential issues associated with maturing unfalsified control theory into a practical technology for application to production heat pumps are discussed.**

I. Introduction

Modern heat pumps are engineered to provide energy efficient cooling and/or heating over a broad range of conditions, through the use of variable speed compressors, variable speed fans, electronically actuated expansion valves, and a set of discrete switches and on/off valves. A heat pump control system operates these continuous and discrete actuators to select a mode of operation, to regulate one or more space temperatures (and possibly one or more humidities), and to enforce a set of process variable constraints, all in an energy efficient and comfortable manner.

Because the inputs interact, the control problem is multivariable. Furthermore, modern heat pumps operate over an increasingly large operating envelope over which the dynamics change, making the control problem nonlinear. For example, in a residential building application, this multi-dimensional envelope may be characterized by a range of indoor and outdoor temperatures, a range of sensible and latent heat loads, and a range of compressor frequencies and fan speeds. Finally, heat pumps operate in a large diversity of applications, which introduces significant plant uncertainty into the control problem. Even though the equipment is precisely manufactured in a factory, the installed plant (from a control perspective) 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, which might range from residential buildings to computer server rooms. As such, perhaps the most important (yet overlooked) control system requirement is that it must provide robust performance over a large operating envelope, for a diverse range of applications.

At the same time, any control solution must be simple. The architecture must be understood and trusted by engineering staff with diverse technical expertise. Start-up, shut-down sequences must simple, robust and testable. The solution must be robust with respect to actuation limits, and sensor and actuator failures. In particular, rules for parameter tuning must be clear and robust, and the number of tunable parameters must be minimized. A controller that might provide superior performance in some applications but requires expert tuning will not be accepted into production.

This paper is an attempt to address partially the robustness challenges associated with heat pump control with two tried and true methods: Gain scheduling [1] and unfalsified control [2], [3], [4]. Starting from an existing baseline control system architecture, we make use of a physics-based model of a particular heat pump coupled to a residential house to study how process gains change as a function of operating condition, and find that much of the variability can be compensated using a simple logarithmic gain schedule. This addresses the nonlinearity caused by and associated with many of the internal process variables with a simple schedule that depends on only one parameter. However, this cannot address plant variation caused by application diversity, much of which is not directly measured. For this we introduce an unfalsified compensator for one of the feedback loops, and show how it is effective in removing candidate compensators that fail to satisfy a given performance requirement.

This paper is organized as follows. In Section II we describe the particular heat pump, its model, and a baseline control architecture. In Section III we present results that show process gain varies logarithmically as a function of measured variables, suggesting a simple gain schedule solution. In Section IV we present the unfalsified control compensation, with simulation results. Some conclusions are drawn in Section V.

II. System Description & Model

The heat pump diagrammed in Fig. 1 consists of a conventional outdoor unit connected by pipes to an indoor air handler unit. The outdoor unit includes a variable speed compressor, an outdoor coil, which is a condensing tube-fin heat exchanger (HEX) with a variable speed fan, and electronic expansion valve (EEV). The indoor air handler unit includes an indoor coil, which is an evaporating tube-fin HEX with a variable speed fan. In cooling mode the cycle operates as follows. The compressor compresses the

[1]Mitsubishi Electric Research Laboratories, Cambridge, MA, USA {bortoff}@merl.com, [2] Engineering Department of Aeronautics and Astronautics, Kyoto University Graduate School, Japan, tsuji.kosei.68z@st.kyoto-u.ac.jp

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)

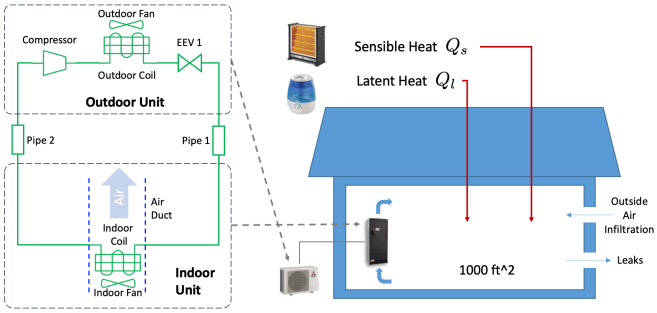

Fig. 1. System diagram.

refrigerant vapor, raising its pressure and temperature, which flows into the outdoor coil, where it condenses into liquid, releasing heat to the outdoor air. The liquid refrigerant enters the EEV which is the expansion device. The EEV orfice size is controlled via a stepper motor, measured in counts. As the refrigerant passes through the EEV, its pressure and temperature drop and it becomes cold, two-phase fluid, which passes through pipe 1, and into the indoor coil, where it entirely evaporates, absorbing heat from the indoor air stream. The low pressure vapor is returned to the compressor via pipe 2, closing the cycle. Heating mode is achieved with a valve which reverses the flow direction, but the basic operation is the same.

A dynamic model of the heat pump was constructed in the Modelica modeling language [5], [6] for purposes of developing control algorithms and evaluating closed-loop system performance in representative buildings and climates. A complete description of the model is beyond our scope, and we refer the reader to [7] and the references. The heat pump system model was coupled to both an adiabatic single-volume room model for understanding the basic dynamics, and a two-zone building model representing a typical residential construction [8], using the Modelica Buildings Library [9].

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). The indoor fan provides ventilation and its speed may be used to regulate zone humidity in some applications, but this is neglected

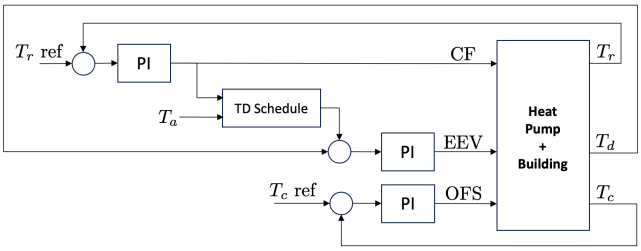

Fig. 2. Baseline control architecture for heat pump control.

in this paper. (It is usually controlled by the indoor unit, whereas CF, EEV and OFS are controlled by the outdoor unit.) 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. If the model is coupled to a building model, additional sensible and latent loads come from the building constructions, weather and air infiltration. If it is coupled to the adiabatic room model, T_a, Q_s and Q_l are the only disturbances. The plant has seven measured outputs: the room temperature and 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.

A. Baseline Control Architecture

Fig. 2 shows a baseline control architecture for the heat pump, operating in either heating or cooling mode. Although the plant is multivariable and interactive, the architecture is decentralized and diagonal, with three feedback loops. The first loop uses CF to control room temperature, using a PI compensator. The second loop uses EEV to regulate T_d to a target value that is computed by a dynamic schedule based on CF and the ambient temperature T_a. This feedback indirectly regulates the evaporator superheat $T_{\mathrm{SH}} = T_s - T_e$ to a small positive value, and is described in [10]. This feedback loop is used instead of one directly regulating T_{SH} because the T_d sensor is more reliable than T_e and it functions better during the start-up sequence, when T_{SH} may be zero. The third feedback loop regulates the condensing temperature T_c to a reference value using the outdoor fan speed OFS. This feedback loop is only active when the outdoor air temperature is sufficiently low, which can reduce the condensing temperature, causing excessive liquid refrigerant to collect in the outdoor coil, starving the indoor coil. In most conditions it is inactive (at a limit) and the OFS is set to a constant value, which may depend on other variables. Each PI includes anti-windup logic to enforce constraints on the actuators. Additional selector logic is used to enforce constraints on other process variables but is beyond our scope.

PI compensators were designed for each loop individually, one at a time, at first ignoring interaction. Disk margins [12] were computed for each loop, one at a time. For example, Fig. 3 shows a Nyquist plot and disk margins for skew $\sigma = 0$ and $\sigma = 1$ for the CF $- T_r$ feedback loop, with the other two loops open, at a nominal condition. Then the

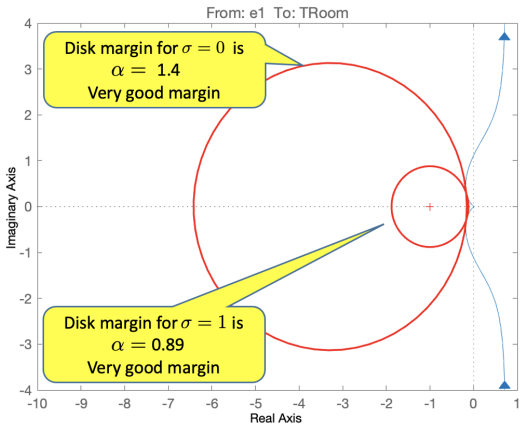

Fig. 3. Disk margin for CF-T_r feedback loop at design condition.

TABLE II
DISK MARGINS ($\sigma = 0$) FOR COMBINATIONS OF ACTIVE FEEDBACK.

LEV	CF	OFS	Margin
T_d	T_r	T_c	1.14
OFF	T_r	T_c	1.34
T_d	OFF	T_c	1.44
T_d	T_r	OFF	1.18
OFF	OFF	T_c	1.65
OFF	T_r	OFF	1.42
T_d	OFF	OFF	1.76

coupled multivariable system was analyzed for robustness by computing multivariable disk margins for all seven 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, because of actuator saturation, or simply because an output variable is at its limit. For example, disk margin values are provided in Table II for cooling mode at a standard design condition. (Note that in a production controller, several additional feedback loops may be selected by logic to enforce a number of output constraints. Margins for all these were also computed and verified similarly, but are omitted for space reasons.)

III. Process Gain Variation & Gain Scheduling

To analyze the variation in plant gain across the operating envelope, it is desirable to first parameterize the operating envelope using appropriate, preferably directly measured variables (to facilitate gain scheduling). Toward that end, a simulation model was constructed in Modelica with three feedback loops closed around the open-loop heat pump model, as shown in Fig. 4. The integral action in the PI feedback effectively computes steady-state values of the two disturbances Q_s and Q_l, which are not directly measured in operation, given values for the measured outputs T_r and ϕ_r. In addition, values of the EEV are computed by feedback to give a desired value of superheat T_{SH}. This arrangement computes steady-state values of the four control inputs and three disturbance inputs to give a desired zone temperature T_r, zone humidity ϕ_r, superheat T_{SH}, for a given ambient

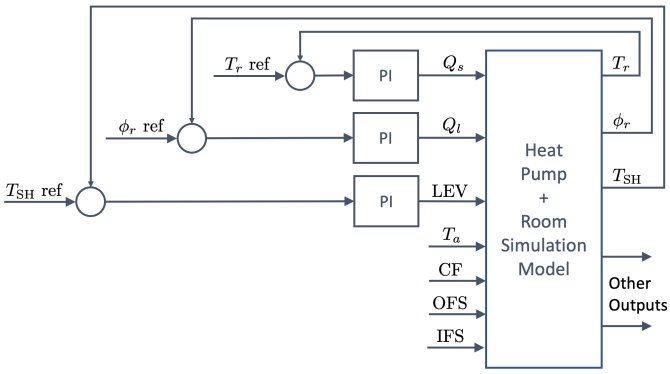

Fig. 4. Feedback control used to compute steady-state plant inputs over operating range.

temperature T_a, compressor frequency CF, and values of indoor and outdoor fan speed IFS and OFS respectively. This allows for parameterization of a seven dimensional operating envelope by the seven measured variables CF, IFS, OFS, T_a, T_r, ϕ_r and T_{SH}. A range of values for each of these variables was used to sample the operating envelope at a total of 840 conditions, and compute steady-state equilibrium solutions for the heat pump model. These were then used with the open-loop heat pump model to compute linearizations for each condition.

Frequency responses were computed from each linearization. Fig. 5 shows two representative plots. At the top of Fig. 5, the frequency response for the CF-T_r pair in the control architecture from Fig. 2 is shown, parameterized by CF. Two features are apparent. First, the low frequency gain changes by 10x over its range from CF=20Hz to CF=120Hz. This large range cannot be compensated effectively by a fixed gain PI compensator. Secondly, the bandwidth is slightly wider at higher frequency. This is minor and less important. The same observations can be garnered from the bottom plot for the OFS-T_c pairing. In fact, it is known from experience that both feedback loops loose margin and can become unstable at low values of CF and OFS, respectively, because the plant gain increases significantly at these conditions.

Closer inspection of the low frequency gain variation reveals that both vary logarithmically as a function of their respective actuator value. Fig. 6 plots the *inverse* of the low frequency gain versus CF (top) and OFS (bottom), showing the measured data fits well with a straight-line model. A linear relation between the actuator variable and the inverse of the low frequency gain implies a linear relation between the actuator variable and the logarithm of the low frequency gain. This suggests a simple gain scheduled PI compensator that depends on only one additional parameter:

$$\mathrm{CF}(t) = \sigma \cdot \left(k_p T_r(t) + k_i \int_0^t T_r(\tau)\, d\tau \right) \qquad (1)$$

where scheduling parameter σ is defined as

$$\sigma = (1 + s \cdot s_n \cdot (\mathrm{CF} - \mathrm{CF}_n)), \qquad (2)$$

where s is the slope of the curve fit in Fig. 6, s_n is the inverse

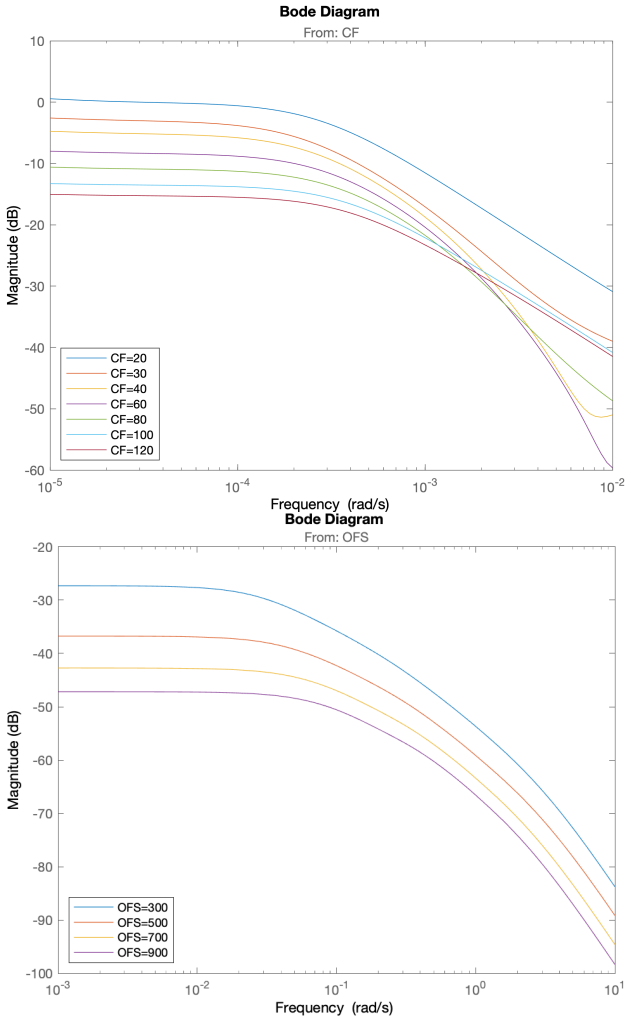

Fig. 5. Frequency responses for CF-T_r (top) and OFS-T_c (bottom) pairs, for a range of steady-state values of CF (top) and OFS (bottom).

of the low frequency gain at a fixed, nominal condition, and CF_n is the nominal condition, e.g. $\mathrm{CF}_n = 40$Hz. Presuming a nominal PI design is done at a nominal condition, s_n and CF_n, k_p and k_i are known, so the gain schedule σ is dependent on only one additional parameter s. A similar schedule may be expressed for the OFS-T_c feedback loop. The EEV-T_d feedback loop may be compensated using a flow model of the EEV and in a similar fashion, but details are omitted.

An analysis of the 840 sets of data determined that most of the process gain variation in the CF and OFS feedback loops was due to variation in CF and OFS themselves. Other variation is relatively small and does not adversely affect robustness. Therefore, these PI compensators were gain scheduled, and the multivariable disk margins described in Section II-A were computed for all 840 conditions, with the result being that all possible combinations of feedback over all tested conditions provide acceptable disk margin. The design was repeated in heating mode using a different set of conditions but the same general process.

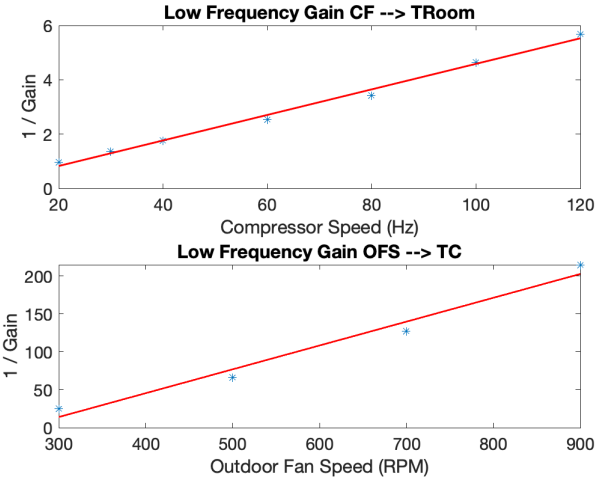

Fig. 6. Linear relation between CF and the inverse of the low frequency gain of the transfer function between CF and T_r (top), and between OFS and the inverse of the low frequency gain of the transfer function between OFS and T_c (bottom).

IV. Unfalsified Control

Gain scheduling is effective to compensate process gain variation over a heat pump operating envelope. However, it is not effective in compensating uncertainty caused by variations in heat pump application and installation. In a typical heat pump installation, variation in pipe lengths and height differences between the indoor and outdoor units can affect the plant dynamics. Thermodynamic coupling between the indoor unit and the building is also a cause of uncertainty. Small versus large rooms, and coupling among rooms can cause a large variation in plant behavior. It is not desirable to manually tune parameters to compensate for these effects at installation time.

Unfalsified control theory [2], [3], [4] offers a promising approach. In unfalsified control, measured plant input and output signals $u(t)$ and $y(t)$, respectively, are used to compute a cost function $J(t)$ for each member of a finite set of candidate compensators. If the computed cost for a candidate compensator exceeds a threshold, it is falsefied by the data, meaning it cannot meet the performance criteria if it were inserted into the feedback loop. It is eliminated from consideration, and if it happens to be in the feedback loop at that time, the controller switches to another candidate that remains unfalsified. The key idea is that a candidate compensator need not actively be in the feedback loop to be evaluated, under mild assumption that it is invertible (such as a PI compensator).

Fig. 7 shows a block diagram of a generic feedback loop with mixed sensitivity performance weights W_p and W_u, and compensator candidates K_i, $1 \leq i \leq 7$. Assume u and y are measured and K_j is in the feedback loop while the other K_i $i \neq j$ are not. Then if each K_i is invertible, a fictitious error $\hat{e}_i$ us computed as $\hat{e}_i = K_i^{-1} * u$, where $*$ denotes the convolution operator, and from this, a fictitious reference $\hat{r}_i = \hat{e}_i + y$ is computed. Then a mixed sensitivity cost such

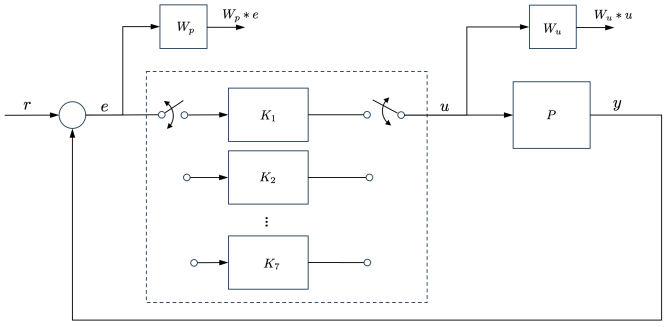

Fig. 7. Block diagram for unfalsified control.

as

$$J_i(r, y, u, t) = \max_{\tau \le t} \frac{\|W_p * e\|_\tau^2 + \|W_u * u\|_\tau^2 - \sigma^2 \tau}{\|\hat{r}\|_\tau^2 + \epsilon} \tag{3}$$

is computed for each candidate compensator. If J_i exceeds a threshold J_{max} at any $t > 0$ then compensator i is falsified, and removed from consideration. It is switched out of the feedback loop if it is active, or removed from future consideration if it is yet to be falsified but not in the feedback loop. Importantly, a compensator can be evaluated without being active in the feedback loop.

We propose this approach for the PI compensator in the $\mathrm{CF} - T_r$ feedback loop. The uncertainty due to installation largely appears to be associated with this feedback loop, since it includes the room temperature and therefore its (uncertain) dynamics, and since this loop is also affected by pipe length variation. The other two loop are less affected by installation variation.

For the design, we choose mixed sensitivity weights

$$W_p(s) = \frac{(s/M + \omega_b)}{(s + A \cdot \omega_b)} \tag{4}$$

with $M = 1.5$, $A = 0.001$ and $\omega_b = 0.001$ and

$$W_u(s) = \frac{0.1s}{(s + 0.1 \cdot \omega_b)}, \tag{5}$$

where ω_b is selected to be near the desired bandwidth, A is selected to give a tight fit, reducing peaks in the sensitivity function, and the form of W_u is chosen to avoid penalty on steady state values of $u = \mathrm{CF}$ that occur in disturbance rejection.

It is straightforward to use the nominal transfer function from CF to T_r from our model and compute the H_∞ mixed sensitivity controller for (4) and (5) using the Matlab `mixsyn` command, giving a cost $\gamma = 1.3$. Then we can compute PI gains that closely approximate this controller, using a formal model reduction or just by inspection of the magnitude plot, shown in Fig. 8. From this we can compute a set of seven PI controllers, with gains ranging from $(1/8)$x to 8x this nominal PI compensator, spaced logarithmically. Frequency responses of the set are shown in Fig. 8, while closed-loop step responses for each compensator are shown in Fig. 9. This shows a span of responses ranging from

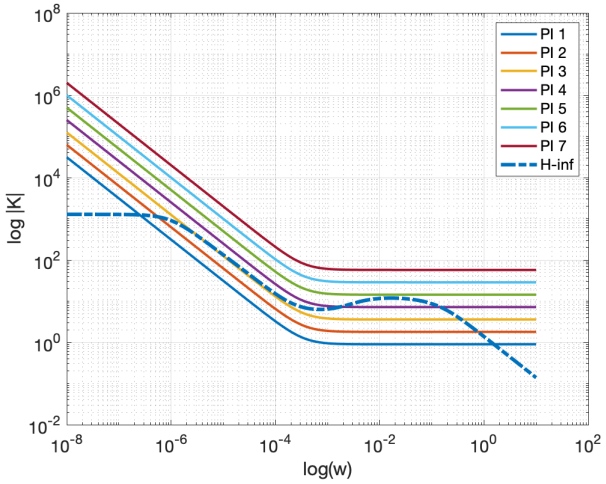

Fig. 8. Frequency response magnitude for a set of seven PI candidate compensators, and the compensator that minimizes the H_∞ mixed sensitivity problem for the CF - T_r feedback loop.

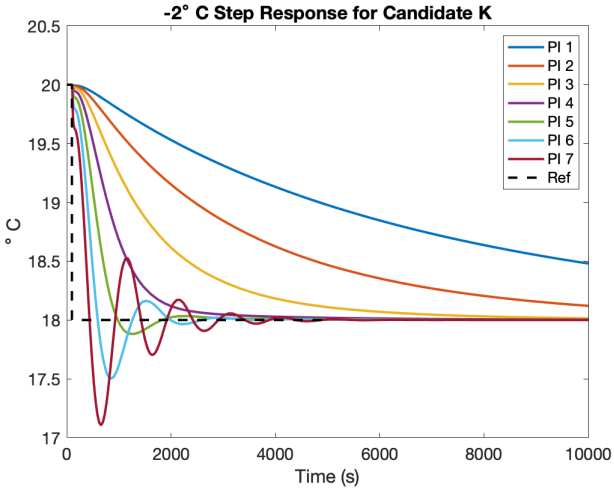

Fig. 9. Closed loop step responses for the set of candidate PI compensators, ranging from high gain and highly oscillatory (PI 1) to low gain and sluggish (PI 7).

oscillatory and nearly unstable to sluggish, with both limits being unacceptable for the nominal case. (However, they may be acceptable in some installations where the plant is different.)

Fig. 10 shows a step response (2°C step down) of the full unfalsified control feedback with the nonlinear Modelica model, including the gain schedules described in Section III, and with nominal PI gains designed using conventional SISO loopshaping techniques. A simple logic is realized to select from unfalsified control candidates: The simulation starts with the highest gain PI ($i = 7$), and almost immediately falsifies $i = 7$ and $i = 6$ because of the high initial cost, due largely to the large value of CF. The logic simply switches to the largest i still unfalsified. After some time, the 3 low-gain sluggish compensators are falsified, leaving 2 unfalsified compensators $i = 4, 5$, which are nearly optimal. For this

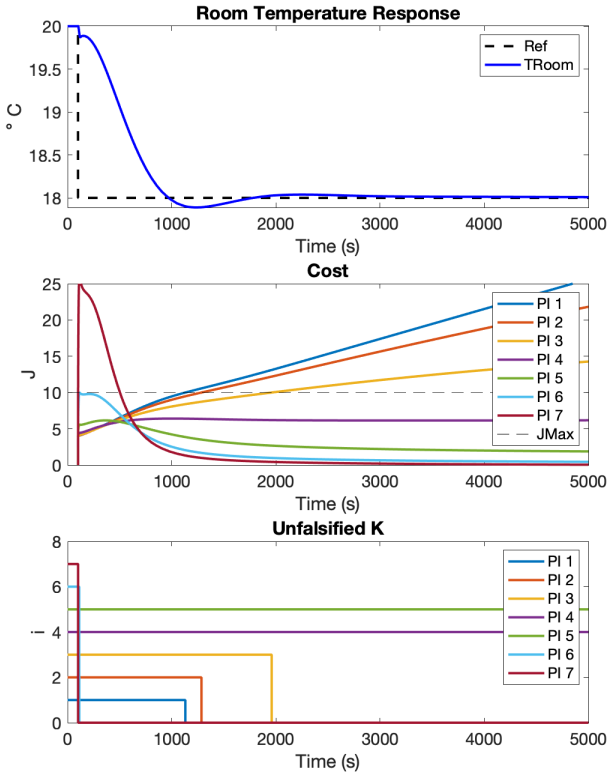

Fig. 10. Unfalsified control step response (top), cost functions for each compensator candidate (middle), and falsification status, with 0 being falsified (bottom).

simulation, $\sigma = 1$, $\epsilon = 1$, and $J_{max} = 10$.

It is interesting to study the cost functions in Fig. 10. Those PIs with high gain are quickly falsified, which is important because these could cause poor performance or instability. The sluggish compensators are eliminated after more time. The best compensators are those that balance both the fast and slow, and this may suggest potentially new criteria for judging performance.

The simulation shows that the unfalsified concept is functional and will eliminate from consideration those controllers who fail to meet the defined performance function. What still remains is a more thorough analysis for a range of plant models that may be constructed by varying those aspects that can be modeled, such as room size and pipe length. The effect of disturbances also requires some additional research.

V. Conclusions

This paper presented a model-based analysis of process gain variability for a residential heat pump application, and identified a relatively simple approach to gain scheduling the main feedback loops in a conventional control architecture, requiring only one additional parameter to be tuned. It also investigated the application of unfalsified control as a potential means to compensate uncertainty associated with variations associated with heat pump installation and application. Both these methods show promise but several questions and challenges remain before these find their way into production.

For the gain scheduled result, it would be interesting to analyze rigorously the physical cause(s) of the logarithmic relation. It seems natural that a linearized plant input, be it CF or OFS or IFS, might be scaled by the value itself, as a kind of percentage normalization. The compressor in this application is a positive displacement type (rotary or scroll) which moves a mass of refrigerant in linear proportion to its speed if boundary conditions are constant, but this is complicated by the coupling with varying pressure and enthalpy at the inlet and outlet.

Unfalsified control is a promising technology because candidate compensators can be evaluated based only on measured data. With the development of remote performance monitoring systems, plant input and output data is increasingly accessible, and may be used on-line or off-line to evaluate new compensators. This is potentially valuable for an engineering function, since the performance of heat pumps has historically been invisible after they leave the factory. On the other hand, there are potential barriers to deploying this into product. The cost function and falsification criteria appears simple, but they contain several parameters that need to be tuned. Automating compensator falsification is risky because it is sensitive to the performance weights, the values of σ and ϵ, and a value for J_{max}. Robust tuning rules need to be established, and issues such as disturbances, time windows and initial conditions all need further investigation.

References

[1] D. Leith and W. E. Leithead, "Survey of gain-scheduling analysis and design," *International Journal of Control*, vol. 73, no. 11, pp. 1001–1025, 2000.

[2] M. G. Safonov and T.-C. Tsao, "The unfalsified control concept and learning," *IEEE Transactions on Automatic Control*, vol. 42, no. 6, pp. 843–847, June 1997.

[3] M. Jun and M. G. Safonov, "Automatic pid tuning: An application of unfalsified control," in *Proceedings of the 1999 Symposium on Computer Aided Control System Design*, 1999, pp. 328–333.

[4] M. G. Safonov, "Data-driven robust control design: Unfalsified control," in *Achieving Successful Robust Integrated Control System Designs for 21st Century Military Applications*, vol. 2, 2006, pp. 1–18.

[5] *Modelica Language Specification Version 3.5*, Modelica Association, https://www.Modelica.org/, Feb. 2021.

[6] P. Fritzon, *Principles of Object Oriented Modeling and Simulation with Modelica 3.3: A Cyber-Physical Approach*. Wiley, 2015.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

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

[12] P. Seiler, A. Packard, and P. Gahinet, "An introduction to disk margins," *IEEE Control Systems Magazine*, pp. 78–95, Oct. 2020.