

Modeling and Coordinated Control of an Air-Source Heat Pump and Hydronic Radiant Heating System

Christopher R. Laughman Scott A. Bortoff Hongtao Qiao

Mitsubishi Electric Research Laboratories, Cambridge, MA, USA
{laughman,bortoff,qiao}@merl.com

Abstract

Hydronic radiant heating systems embedded in building constructions are receiving increased interest due to their potential for high energy efficiency and improved thermal comfort, but their slow time constants pose challenges when controlling space conditions. We address this problem via a system architecture that combines the radiant heating system with a separate air-source heat pump serving the same space. In this paper, we develop a new coordinating control method for this proposed system by using a set of reduced order models generated from a set of coupled Modelica models of the individual subsystems. This new control architecture does not require significant modification of standard heat pump control architectures, and results in both improved thermal comfort and reduced energy consumption.

Keywords: radiant heat transfer, heat pump, thermally active building systems, control, Modelica

1 Introduction

Buildings continue to be a focused target of societal efforts to reduce carbon-based energy consumption, as they consume a significant fraction of the total energy produced in many countries. Improving energy efficiency of equipment, systems and buildings has been a long-term trend since the mid 1970s. A more recent trend is the accelerated growth of renewable energy generation, which makes electric heat pumps, for example, an increasingly attractive alternative to conventional boilers and furnaces. In addition, increased awareness of the impact that buildings have on the health and safety of their occupants is prompting extensive study of approaches that improve the overall thermal comfort of occupied spaces, rather than just air temperature.

As described in Rhee et al. (2017), this set of concerns has motivated the increased exploration and application of radiant heating and cooling systems. The main benefit of these systems is that they are able to provide equivalent thermal comfort at a lower temperature difference, due to their extended surface areas and the fact that the

mean radiant temperature experienced by an occupant contributes significantly to thermal comfort. The smaller difference between the desired room temperature and the source temperature generally makes it possible for heating and cooling systems to operate with higher energy efficiency. These systems also tend to be quiet, as they do not rely upon the transport of air to remove the thermal load, and can be integrated into building systems, such as when the polyethylene tubes transporting the water most often used as a heat transfer medium are embedded in structural concrete slabs used in the building. Such radiant heating and cooling systems that are integrated into the building envelope are often called thermally-active building systems (TABS). The high thermal inertia of TABS is beneficial because it can facilitate the reduction of energy consumption peaks by allowing the space to be conditioned when occupants are absent and the cost of electrical power is low. The corresponding reduction in the required heating and cooling capacity also can reduce the size and cost of system components.

While TABS have a number of advantageous characteristics, the significant thermal inertia associated with the time constants of the slab and low temperature differences in the system impose constraints on their controllability (Gwerder et al., 2008). In particular, the long time delays in the thermal response of the slab, which typically range from hours to days, make it either impractical or impossible to change the slab temperature in response to sudden changes in the room air temperature setpoint or large changes in the room load. Consequently, the dynamic response due to heat load disturbances in a room conditioned by a TABS is much larger than it would be if the room were conditioned by a heat pump or other more conventional system. The slow TABS system lacks the control authority and bandwidth to reject the heat load disturbance.

A variety of previous work has been done to address these challenges in applying radiant heating systems. The control of on/off valves have been an object of particular study due to their low cost; for example, Tang et al. (2018) developed a 3-D discretized model of a TABS to describe the transient thermal dynamics of the slab to study a pulsed flow control strategy under cooling




Figure 1. Diagram illustrating building model incorporating heat pump and thermally-active slab for radiant heating.

operation while assuming constant thermal conditions on the slab boundaries. They determine through simulation and experiment that the water flow rate can be reduced by approximately 25% while modulating capacity over a wide range of values. Gwerder et al. (2009) also explores the control of these systems using pulse width modulation strategies with a limited number of tuning parameters, and demonstrates these methods' efficacy in both simulation and experiment. Márquez et al. (2017) investigated a state machine-based control strategy for controlling the heating operation of a TABS, a fan coil unit, or a combination of both systems by using a TRN-SYS model of the space and a curve-based model of a heat pump to provide the water supply, and demonstrate that the thermal comfort and energy consumption are strongly dependent upon the control methods used. Beghi et al. (2011) also investigated a relay controller architecture that manages the water flow rate through both radiant slabs and fan coils using a dynamic model of the space and curve-based models of the fan coil to directly control thermal comfort during both heating and cooling operation, and reduced energy consumption by 5-13% via this approach. Finally, Gayeski et al. (2011) used a series of temperature- and load-dependent maps for a variable capacity chiller to perform predictive precooling of a concrete-core radiant floor, and experimentally demonstrated energy savings of 19-25% in Atlanta and Phoenix climate conditions.

One system architecture that addresses the limitations on the control bandwidth of the TABS, which are fundamental to its construction, involves the combination of a TABS and a variable-capacity air-source heat pump, as seen in Figure 1. This combination of systems provides a means to reject thermal disturbances and provide rapid setpoint changes while also providing the increased thermal comfort and efficiency of the TABS. While such an architecture has a variety of benefits, one inherent challenge is that both the heat pump and the TABS will have a

high control gain to the temperature and thermal comfort of the room. The operation of the parallel systems must therefore be coordinated to simultaneously minimize the overall energy consumption and maintain thermal comfort.

The dynamic nature of the interaction between these subsystems necessitates the use of dynamic models and controls to identify a feasible design. Because model-based control design practices can be an invaluable tool to quickly and inexpensively explore the control design space to develop and refine candidate controller concepts for nonlinear thermofluid systems, we use a set of dynamic models, as opposed to performance maps, of both the air-source heat pump and the room with the TABS to design the controls for the overall system. As has been noted in Wetter et al. (2016), the equation-oriented language Modelica (Modelica Association, 2017) is well-suited for such system-level analysis and design tasks because of its ability to separate concerns between the activities of physics-based modeling and simulation, and because it facilitates the use of models beyond simulation, such as numerical optimization or linearization for control design. Such capabilities make it possible to make individual dynamic models of the air-source heat pump, the TABS, and the room-air dynamics, and then to couple all of these models together in a heterogeneous simulation for the use in the design of the coordinated control system and subsequent nonlinear simulations.

In this paper, we describe a set of Modelica models that represent the coupled dynamics of an air-source heat pump and a room with a TABS during heating operation, and then describe a candidate approach for coordinating the operation of both systems to meet thermal comfort guidelines while managing the power consumption. Our objective is the design of a control architecture that requires minimal modification of existing heat pump control algorithms, while providing a high degree of energy efficiency and thermal comfort. A secondary objective is the illustration of the use of Modelica models for frequency-domain control system design, which is a use case that is different than time-domain simulation. In Section 2, we describe the models used, first for the air-source heat pump operating in a heating configuration, and then for the thermally-active slab. We also describe a few salient details about the Modelica implementation as well as a number of adaptations required to make the individual system models compatible. We then linearize the nonlinear model and develop a control design for the overall system in Section 3 that meets a specified set of requirements and objectives, and evaluate the performance of this controller in simulation. Finally, we briefly review a set of conclusions and next steps in Section 5.

2 Models

A control-oriented analysis of the dynamics of the proposed heating systems requires a two distinct models: a model of the vapor compression heat pump, and a model of the room with the TABS. As the dynamics of the integrated system are of primary concern, the subsystem models must correctly capture their behavior and encode it in a representation that can be coupled with the other system models. This analysis thus builds on previous modeling work relating to the subsystems; these heat pump models and information about correlations and further details are described at length by Laughman et al. (2015), while the building models are similarly discussed in Wetter et al. (2014). We will briefly review the structure of these subsystem models and then describe the manner by which they were integrated into an overall system model in the following subsections.

2.1 Vapor Compression Heat Pump Model

A conventional variable-capacity air-source vapor compression cycle was used as the basis for the development of this model. This cycle includes a variable speed rotary compressor, a refrigerant-to-air condensing heat exchanger with a variable-speed fan in the occupied space, an evaporating heat exchanger with a variable-speed fan in the ambient environment, and an electronic expansion valve. The availability of four control inputs provides the capability to regulate both the internal refrigerant-side variables of the cycle, such as pressures or condenser subcooling, as well as overall cycle performance variables, such as its heating capacity or outlet air temperature. We also assumed the availability of a variety of temperature measurements for the purposes of control, but did not assume that other measurements were available that could not be realistically included in a commercial product, such as mass flow meters.

In considering the behavior of the cycle, the dynamics of the compressor and electronic expansion valve are much faster than those of the heat exchangers, so that the heat exchanger dynamics dominate the system response over the time scales of interest (multiple seconds to hours). We consequently used static (algebraic) models for the compressor and expansion valve, and dynamic models of the heat exchangers. A finite volume discretization approach was used to describe the behavior of both heat exchangers to ensure the physically correct treatment of mass, momentum, and energy conservation. A number of simplifying assumptions were used in these models; the most important of these is that of one-dimensional refrigerant flow, meaning that the flow field is assumed to be uniform in the r and θ directions at each point along the length of the heat exchanger. Other important assumptions included that of a Newtonian fluid, negligible axial heat conduction along the direction of the fluid flow, negligible viscous dissipation in the fluid,




Figure 2. Staggered discretization grid used to model the dynamics of the fluid.

negligible contributions to the energy equation from the kinetic and potential energy of the refrigerant, negligible dynamic pressure waves in the momentum equation, and thermodynamic equilibrium in each volume for which the refrigerant is in the two-phase region.

The resulting conservation equations are then discretized using the upwind approximation and can be written as

$$\frac{d(\rho_j V_j)}{dt} = \dot{m}_k - \dot{m}_{k+1} \quad (1)$$

$$\frac{d(\dot{m}_i)}{dt} l = \rho_j v_j^2 A_j - \rho_{j+1} v_{j+1}^2 A_{j+1} + \frac{A_j + A_{j+1}}{2} (P_{j+1} - P_j) + F_{f,i} \quad (2)$$

$$\frac{\partial(\rho_j u_j A_j)}{\partial \tau} = H_k - H_{k+1} + v_j A_j (P_{j+1} - P_j) + v F_{f,i} + Q_j. \quad (3)$$

These equations are implemented on a staggered flow grid to eliminate numerical oscillations in the pressure, which is illustrated in Figure 2. The indices in Equations 1 through 3 match those used in the figure, where the i indices refer to the momentum grid, the j indices refer to the thermal grid, and the $k = j + 1$ indices refer to the boundaries of the thermal grid for each volume V_j . In addition, the term H_k is defined as

$$H_k = \dot{m}_k \bar{h}_{upstream,j}, \quad (4)$$

and the mixed-cup specific enthalpy \bar{h} is equal to the *in situ* specific enthalpy under the homogeneous flow assumption.

A set of simplified closure relations for the frictional pressure drop related to F_f and the refrigerant-side heat transfer coefficients were used to connect those variables to the properties of the refrigerant flow through each control volume. Simplified forms were used because many published correlations from the literature have poor numerical properties that make them very difficult to include in a dynamic simulation. The frictional pressure drop was expressed as

$$\Delta P = K \frac{(\Delta P)_0}{\dot{m}_0^2} \dot{m}^2, \quad (5)$$

where the nominal values of K , $(\Delta P)_0$, and \dot{m}_0 were determined by using the Colebrook correlation for the single-phase friction factor and the Friedel correlation for two-phase multipliers. Similarly, phase dependent heat transfer coefficients were used, in which the value of the heat transfer coefficient in each phase was only dependent on the refrigerant mass flow rate, and trigonometric interpolation was used to smooth the transition between phases over a small range of thermodynamic quality (Richter, 2008). The constants used for this simplified correlation were calculated by the Gnielinski correlation for single-phase fluids, the Shah correlation for condensing flows, and the Gungor-Winterton correlation for evaporating flows.

The pressure P and mixture specific enthalpy h in each control volume were used to define the thermodynamic state of the refrigerant, as well as the state derivatives used in the numerical integration routines used to simulate the behavior of the cycle. The derivatives of the mass and energy in the control volume were formulated in terms of these state variables. While this selection of state variables can result in fluctuations in total refrigerant mass in the cycle due to accumulated numerical errors (Laughman and Qiao, 2017), these errors were found to be negligible in this work.

A moist-air formulation was used to describe the heat transfer from the outer surfaces of the tubes to the air, as described in Equation 7, where the mass transfer coefficient was given by a modified Lewis correlation. The humidity dynamics were negligible in these models due to heating mode operation, but were included for compatibility with the room models.

$$\dot{m}_{air} C_{p,air} \frac{dT_{air}}{dy} \Delta y = \alpha_{air} (A_{o,tube} + \eta_{fin} A_{o,fin}) (T_w - T_{air}) \quad (6)$$

$$\dot{m}_{air} \frac{d\omega_{air}}{dy} \Delta y = \alpha_m (A_{o,tube} + \eta_{fin} A_{o,fin}) \times \min(0, \omega_{water,sat} - \omega_{air}) \quad (7)$$

The cycle model also included a variable-speed high-side rotary compressor, in which the motor is cooled by the compressed high-pressure refrigerant discharged from the compressor mechanism. While mass and energy balance equations can be formulated for this system, a model of realistic machine performance must take into account aspects of the machine behavior that are difficult to capture with low-dimensional models. The compressor was therefore described via the volumetric efficiency η_v and isentropic efficiency η_{is} , as well as a function for the power consumption $\dot{W}(P_{rat}, \omega)$ that relates the total power consumption of the machine to other refrigerant-

Parameter	Value
Refrigerant	R410A
Total refrigerant mass (kg)	0.529
OU HEX tube diameter (mm)	9.5
IU HEX tube diameter (mm)	7.9
OU HEX tube length (m)	1.1
IU HEX tube length (m)	0.9
OU HEX number of tubes	26
IU HEX number of tubes	18

Table 1. Cycle model parameters.

side parameters, given by

$$\eta_v = \frac{\dot{m}_{comp}}{\rho_{suc} V f} \quad (8)$$

$$\eta_{is} = \frac{h_{dis,isen} - h_{suc}}{h_{dis} - h_{suc}}, \quad (9)$$

and the formulation of the expressions for η_v , η_{is} , and \dot{W} are provided in Laughman et al. (2017).

A simple isenthalpic model of the electronic expansion valve was also used, as described by a standard orifice flow equation

$$\dot{m} = C_v a_v \sqrt{\rho_{in} \Delta P}, \quad (10)$$

where the flow coefficient C_v can be determined by regression using experimental data and the flow area a_v is specified by the user. Algebraic fan models specified by standard fan laws (ASHRAE, 2008) were also used, and were scaled by experimentally measured values of the fan speed, flow rate, and power for a representative system. The building was also assumed to maintain a constant pressure, allowing us to assume that the supply air flow rate was equal the exhaust air flow rate in the conditioned space.

Following the construction of these component models, they were interconnected to form a complete cycle model. This model was configured with a number of geometric and fluid parameters that governed the behavior of the overall system, which are given in Table 1. The resulting system had a nominal heating capacity of 2631 W at a compressor frequency of 60 Hz, with 2.5 °C of condenser subcooling when the indoor zone was at 26 °C and 2.3 °C of evaporator superheating when the ambient environment was at 2 °C.

2.2 Building Model

The increasing adoption of the Modelica language by many different engineering communities has resulted in the growth of an ecosystem of high-quality open source libraries that can be freely used and incorporated into other models. The Modelica Buildings library (Wetter et al., 2014) is a prime example of such a library, as it has been under development for 10 years with the extensive support of many researchers and organizations and has

Parameter	Value
Room area (m ²)	55.7
Room height (m)	3.66
Concrete slab thickness (m)	0.33
Window size (m ²)	13.9
Building location	Boston, MA, USA

Table 2. Room model parameters.

been extensively tested and used. This library is designed to provide dynamic models of moist air and other media, heat transfer, multizone airflow, and electrical distribution for building and district energy and controls methods; the many component models that are provided in the library can be assembled into models of rooms and single- or multi-story buildings. We employed this library to simulate the building physics in this work; a diagram of this model is illustrated in Figure 1, with pertinent parameters provided in Table 2.

The multiphysical and object-oriented nature of this modeling paradigm is reflected in the structure of the room models developed for this work, as the constructions that constitute the room are composed of fundamental materials and elements that are directly related to the building materials used in practice. Models of the walls and the room are composed of multiple layers representing the individual materials, each of which is parameterized by fundamental material properties like thickness, thermal conductivity, and density. These layers are further discretized into control volumes to approximate the behavior of the partial differential equations describing heat conduction in the material. The default number of control volumes used for a layer of material is automatically chosen, based upon the Fourier number of the material, so that the time constants of each volume are approximately equal. One wall of the room model also incorporates a double-paned window construction over a large portion of its surface. Further information about these models is provided in Wetter (2006).

A similar level of detail was employed to construct the physical models of the zone air and the thermal loads. The zone air model is a mixed air single-node model with one bulk air temperature that interacts with all of the radiative surfaces and thermal loads in the room, where the zone is assumed to have convective, radiative, and latent gains specified on a per area basis. The radiative heat transfer representing the solar heat gains and the infrared heat transfer between the interior surfaces of the room is also modeled with a similar degree of accuracy. The absorptivity and emissivity of each surface is characterized, and is used in the calculation absorbed and reflected heat transfer to all of the surfaces in the room (Wetter et al., 2011). A set of simplified view factors between the surfaces in the room was used as an approximation to avoid the complexity of incorporating the detailed room geometries. Detailed weather data from TMY3 files was also used to describe the ambient conditions, including

but not limited to the dry bulb temperature, wet bulb temperature, and direct and diffuse solar radiation terms, to accurately describe the influence of these conditions on the room thermal dynamics.

The overall building model was defined by using the standard interface established in the Buildings library, in which the room constructions can either be defined as part of the room or only defined via their surfaces, to establish a partial model was developed that only defined the surface absorptivity and emissivity properties of the slab. This partial model could thus be extended to incorporate either an inactive slab or a thermally active slab. The inactive slab used a standard construction with a 8 cm layer of concrete, a 5 cm layer of insulation, and a second 20 cm layer of reinforced concrete, whereas the thermally-active slab used a similar structure but also incorporated crosslinked polyethylene (PEX) tubes to provide a radiant thermal input. The overall model for this slab is based on models described in Koschenz and Dorer (1999), in which the PEX tubes embedded in the concrete were modeled using a simplified fluid model, and a limited (5) number of regions were used to discretize the pipe to limit the computational complexity of the overall system. The inclusion of the TABS in our system provided two additional control inputs from which to influence the thermal comfort of the room: the mass flow rate of water through the TABS, and the inlet temperature of that water.

Conventional physical boundaries between subsystems were used as abstraction layers for these models. The cycle model of the heat pump was encapsulated as one discrete model, where the flow and temperature sources for the outdoor unit were wrapped into the overall cycle model and two air-side connectors constituted the main interface between the cycle model and the room model. Similarly, the room model was also encapsulated as a separate discrete model with air-side connectors to interface with the heat pump. This overall system model incorporating the connected room and cycle models was then itself wrapped in another enclosing model that exposed the limited set of inputs and outputs used for the controller design through the generation of linearized system models in Dymola, and performance assessment through the testing of the controller performance on a single plant model.

One technical concern that emerged in the course of this work is that the standard moist air media model used in the Modelica Standard Library (MSL) and the simplified moist air model provided in the Buildings Library are not interchangeable, due to the different assumptions made in these models. While `Modelica.Media.Air.MoistAir` uses a mixture model in which both the air and steam components obey the ideal gas law $P = \rho RT$, `Buildings.Media.Air` simplifies this model further so that the pressure and tempera-

ture are decoupled via

$$\rho = \rho_{STP} \frac{P}{P_{STP}}, \quad (11)$$

where STP denotes the moist air properties at a temperature of 273.15 K, a pressure of 101325 Pa, and a humidity ratio of 10%. This simplification is used in the building models because the nonlinear equation blocks in the flattened model are much smaller, resulting in an increase in simulation speed as well as a marginal decrease in accuracy. Because the air-side models of the heat exchangers for the vapor compression cycle did not assume this simplification, we initially used the MSL media model in the building model, but found that the simulation was very slow. Rather than modify the air-side model of the heat exchangers, we slightly modified `Buildings.Fluid.Interfaces.ConservationEquation` so that the mass in the control volume m was dependent only on the pressure in the volume, e.g.,

$$m = \text{fluidVolume} * \text{medium.p} / P_{STP} * \rho_{STP}, \quad (12)$$

where `fluidVolume` represents the volume of the room.

3 Control Design

The set of system models described in Section 2 can be used to design a control system that coordinates the operation of the heat pump and the TABS. This design is based on frequency-domain methods and demonstrates an important use case of Modelica models that goes beyond time-domain simulation.

Contemporary variable-capacity heat pumps have four control inputs that are used to regulate their behavior, including the compressor speed f_c , an electronic expansion valve (EEV) position, and the indoor and outdoor fan speeds. We assume that the heat pump is controlled using two PID feedback loops: one that actuates f_c based on room air temperature T_a , and a second that actuates EEV based on an internal measurement y_1 such as condenser subcooling. In comparison, the indoor fan speed is typically controlled by the customer, and the outdoor fan speed is used for purposes outside our scope. Because this and other similar control laws are well-established in commercial heat pumps, we do not wish to modify this architecture in the coordinated controller. Furthermore, we assume that the inlet water temperature T_w is the available control input for the TABS.

3.1 Linearization and Model Reduction

Our control design process begins by computing a linearization of the overall system that includes the heat pump, the TABS and the building at a nominal operating condition. The models described the previous section were coupled to refrigerant property models from

the Vapor Cycle Library Modelon AB (2018), compiled in Dymola 2018 FD01 AB, and run on a PC with an Intel i7 processor with 16 Gb of RAM. We simulate the model with constant inputs for 5×10^6 s (58 days) to ensure that the linearization is computed an equilibrium solution. This ensures convergence of the slowest mode of the system, which has a time constant of approximately 1 week. For our particular system the resulting raw linearized model has 344 states and is 99.5% sparse.

The raw model is then carefully reduced in dimension. Frequency-based design methods such as loop-shaping (or any formal control design method, for that matter) can fail if the raw model is used directly. This is because the raw model is highly sparse, stiff and numerically ill-conditioned. In fact, many Matlab Control Toolbox commands, e.g., `bode`, reduce the model automatically, with little visibility to the user, and can produce misleading or erroneous results.

Our model reduction process requires several steps. We begin by symbolically eliminating the states associated with refrigerant energy and mass using the variable names that accompany the state vector in the linearized model, as these states are unobservable, uncontrollable and have eigenvalues at the origin. This results in a 240-dimensional model. We then compute a balanced realization of the system and its associated Hankel singular values, which are typically logarithmically spread as a function of frequency, and then truncate the balanced model, keeping approximately 15 states in a partially reduced-order model. As a result of this calculation, we frequently encounter very slow non-physical, weakly controllable and observable modes which we speculate are associated with the closed-loop topology of the vapor compression cycle and airflow in buildings. These non-physical modes are removed by modal decomposition. In addition, we also compute a singular perturbation if the partially reduced-order model includes pressure states. This process typically produces a 10^{th} -order reduced-order model that is numerically well-conditioned and whose spectrum has modal time constants ranging from 20s to 1.5 days, as is expected from practical experience.

3.2 Dynamic Analysis

Figure 3 shows the frequency response of the reduced model from the two actuators, f_c and T_w , to the room air temperature T_a and room mean radiant temperature T_{rad} . We see that f_c has strong authority over T_a over a frequency range 0 – 0.01 rad/s (10 min), but *neither* input is effective at controlling T_{rad} beyond 0.0002 rad/s (8 hours). In fact, both T_w and f_c are equally effective at regulating T_{rad} at frequencies below 0.0002 rad/s. Fortunately, the disturbance bandwidth is correspondingly similar, as can be seen in Figure 4, which shows the frequency responses from the radiative and convective heat loads Q_{rad} and Q_{sen} , respectively, to the temperatures T_a and T_{rad} . This figure demonstrates that both disturbances




Figure 3. Frequency response of T_a and T_{rad} to controls f_c and T_w .




Figure 4. Frequency response of T_a and T_{rad} to disturbances Q_{rad} and Q_{sen} .

have a low-bandwidth effect on T_{rad} , while only the gain from Q_{sen} to T_a has a relatively broad bandwidth. Moreover, both are within the effective bandwidths of the controls.

3.3 Controller Design

Controlling both T_a and T_{rad} is important because perceived comfort, as measured by the predicted mean vote (PMV) (ASHRAE, 2017), weights both T_a and T_{rad} equally for low levels of human metabolic activity associated with typical office conditions. Room occupants are therefore equally sensitive to both T_a and T_{rad} , though only T_a is typically measured. These facts, together with our requirement to use the production heat pump control algorithm, suggests the control architecture shown in Figure 5.

We consider the heat pump and radiant systems as two separate actuators that both contribute to T_a and T_{rad} (not shown). In this architecture, the PID controller is designed to regulate the room air temperature to the set-point r by controlling the compressor frequency of the heat pump f_c while neglecting the radiant system entirely. This PID command is filtered through a low-pass filter (LPF) whose bandwidth is tuned to approximately 4 hours to speed-up the radiant system response. The LPF splits the frequency content of the PID output into a low-frequency component, which is applied to the radiant system, and a high-frequency component that is applied to the heat pump. The compensator C_2 is set equal to P_1/P_2 , where P_1 is a low-order approximation of the transfer function from f_c to T_a , and P_2 is a low-order approximation of the transfer function from T_w to T_a . These approximate models are computed from the frequency

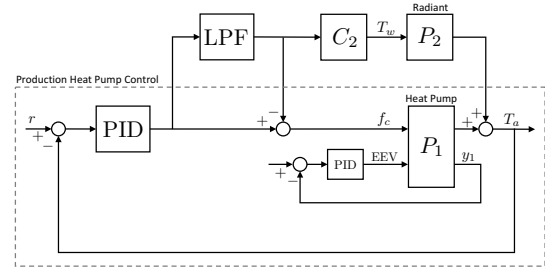


Figure 5. Closed-loop block diagram of the overall system. The production heat pump control is inside the dashed box.

responses described in the previous section and are both minimum-phase, so that C_2 is stable and realizable. We find that a first-order P_2 is sufficient, while P_1 is second-order. The compensator C_2 effectively provides phase lead to the radiant system, so that $C_2 \cdot P_2$ is approximately equal to P_1 , which implies that the frequency response of the compensated radiant system is equal to that of the transfer function from f_c to T_a . The low-pass filter ensures that only low-frequency signals are applied to the compensated radiant system. One benefit of this architecture's simplicity is that the only parameters required by the controller are the low-order models P_1 and P_2 , as well as the gain and bandwidth of the low-pass filter that controls the frequency split and steady-state load split between the heat pump and radiant system. Such an architecture does not require any significant modification of the heat pump controller, aside from subtracting the output of the LPF from the output of the PID, and applying the result to f_c . Note that the complete design depends on only the frequency response.




Figure 6. Room temperature step response. The air temperature reference is stepped 1 °C at $t = 0$, and a heat load disturbance step of $Q_{sen} = -1$ kW is applied at $t = 250$ min. T_a response is similar for both, but the radiant response is sped-up TABS with coordinated control (bottom).

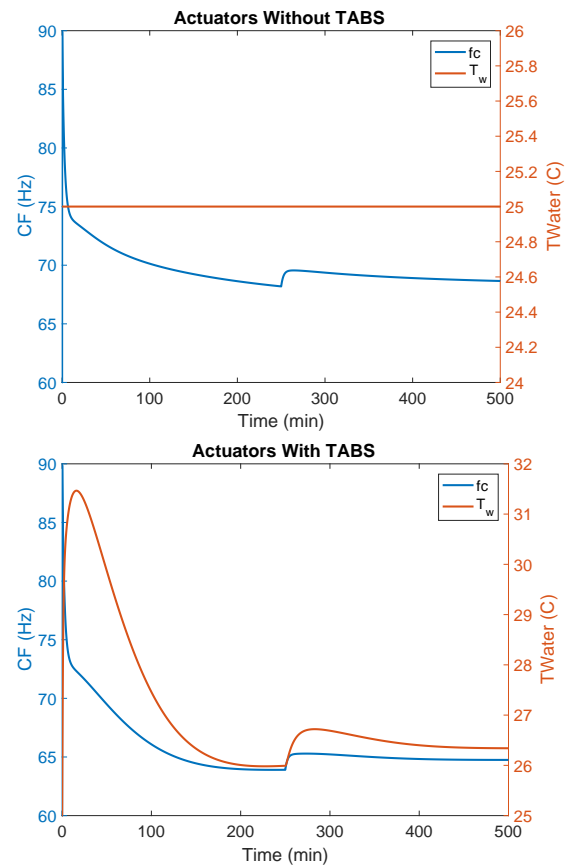


Figure 7. f_c and T_w Actuation. The air temperature reference is stepped 1 °C at $t = 0$, and a heat load disturbance of $Q_{sen} = -1$ kW is applied at $t = 250$ min. f_c is similar for the short term, but f_c is reduced because of the effect of the radiant system (bottom), and in the steady-state, the load is split among the systems. The effect of the lead filter on T_w is apparent and not too large.

4 Results

We first simulate the operation of the system without the TABS. The shortcomings are readily seen by considering the responses of T_a and T_{rad} to a step in the room temperature setpoint (at $t = 0$ s) and to a -1 kW heat load step disturbance at $t = 250$ min, as seen in Figure 6. While the air temperature T_a responds quickly and the disturbance is rejected, the radiant temperature responds much more slowly because of the limited control authority of the heat pump over the surface temperatures of the room. This slow change in T_{rad} implies that a building occupant will feel about 0.5 °C colder for a sustained period of time after the room temperature setpoint change due to the slow radiant temperature response, despite the fact that the air temperature responds quickly.

In comparison, the performance of the heat pump and TABS system under coordinated control is much improved. The response of T_a is nearly identical to that of the previous case, but the response of T_{rad} is significantly faster. While the time constant of T_{rad} for the room without coordinated control is nearly 6 hours, the time constant of T_{rad} is closer to 150 minutes when the

two systems are operated together. This results in a much smaller difference between T_a and T_{rad} , and consequently improved thermal comfort. The disturbance rejection properties of the system are still quite good, as would be expected due to the fact that the heat pump is largely responsible for these dynamics.

The actuator commands for this simulation are shown in Figure 7. Whereas the steady-state compressor frequency f_c is approximately 68 Hz for the base control case, this is reduced by approximately 5 Hz for the coordinated control architecture. This reduction in compressor frequency is also accompanied by a reduction in the needed heating capacity of the heat pump in steady-state for the coordinated control case. As a result, the benefits of this coordinated control method are both a reduction in the power consumption of the heat pump, and potentially a smaller capacity heat pump is required. In addition, consideration of the T_w signal indicates that the improvement in the T_{rad} response is achieved with only a 6 °C increase in the inlet water temperature, which is easily achieved by commercially available water heating systems.




Figure 8. Frequency response of C_s showing lead behavior between 10^{-4} and 10^{-2} rad/s.

5 Conclusions and Future Work

In this paper we have described a Modelica model of a heat pump and radiant heating system integrated into a building, and used the model to construct a coordinating feedback control algorithm using frequency-domain methods. This represents a good use-case example of Modelica models beyond that of time-domain simulation. The use of these linearized models, which are readily computed using tools such as Dymola or OpenModelica, requires some care because they represent large-scale, sparse, and stiff systems. We show a simple coordinated controller for the combined radiant and heat pump system provides improved comfort and responsiveness with a minimum impact on the production control algorithm in the heat pump.

This result can be extended in a variety of different directions. For example, it would be interesting to consider acausal low-pass filters for frequency separation, which may be possible if the references are known apriori. Such an approach constitutes a form of predictive control, and may improve the results and allow further downsizing of the heat pump. Formal MPC methods could also be applied to this same end. Alternate directions could also include the coordination of this system architecture with ventilation systems, as well as consideration of cooling mode operation for these system architectures in light of the effects of humidity.

References

- Dassault Systemes AB. Dymola. 2018 FD01.
- ASHRAE. *HVAC Systems and Equipment Handbook*. ASHRAE, Atlanta, GA, 2008.
- ASHRAE. Thermal environmental conditions for human occupancy. ASHRAE 55:2017, Atlanta, GA, USA, 2017.
- A. Beghi, L. Cecchinato, and M. Rampazzo. Thermal and comfort control for radiant heating/cooling systems. In *IEEE Conference on Control Applications*, 2011.
- N.T. Gayeski, P.R. Armstrong, and L.K. Norford. Predictive pre-cooling of thermo-active building systems with low-lift chillers. *HVAC&R Research*, 18(5):858–873, 2011.
- M. Gwerder, B. Lehmann, J. Tödtli, V. Dorer, and F. Renggli. Control of thermally-activated building systems (TABS). *Applied Energy*, 85:565–581, 2008.
- M. Gwerder, J. Tödtli, B. Lehmann, V. Dorer, W. Güntensperger, and F. Renggli. Control of thermally activated building systems (TABS) in intermittent operation with pulse width modulation. *Applied Energy*, 86:1606–1616, 2009.
- M. Koschenz and V. Dorer. Interaction of an air system with concrete core conditioning. *Energy and Buildings*, 30:139–145, 1999.
- C.R. Laughman and H. Qiao. On the influence of state selection on mass conservation in dynamic vapor compression cycle models. *Mathematical and Computer Modelling of Dynamical Systems*, 23:262–283, 2017.
- C.R. Laughman, H. Qiao, V. Aute, and R. Radermacher. A comparison of transient heat pump cycle models using alternative flow descriptions. *Science and Technology for the Built Environment*, 21(5):666–680, 2015.
- C.R. Laughman, H. Qiao, S.A. Bortoff, and D.J. Burns. Simulation and optimization of integrated air-conditioning and ventilation systems. In *Proc. of the 15th IBPSA Conference*, pages 660–669, 2017.
- A.A. Márquez, J.M.C. López, F.F. Hernández, F.D. Muñoz, and A.C. Andrés. A comparison of heating terminal units: Fan-coil vs. radiant floor, and the combination of both. *Energy and Buildings*, 138:621–629, 2017.
- Modelica Association. Modelica specification, Version 3.4, 2017. URL www.modelica.org.
- Modelon AB. *Vapor Cycle Library User Guide*, 2018. v2.0.
- K-N. Rhee, B.W. Olesen, and K.W. Kim. Ten questions about radiant heating and cooling systems. *Building and Environment*, 112:367–381, 2017.
- C.C. Richter. *Proposal of New Object-Oriented Equation-Based Model Libraries for Thermodynamic Systems*. PhD thesis, Technical University of Braunschweig, 2008.
- H. Tang, P. Raftery, S. Schiavon, J. Woolley, and F.S. Bauman. Performance analysis of pulsed flow control method for radiant slab system. *Building and Environment*, 127:107–119, 2018.
- M. Wetter. Multizone building model for thermal building simulation in Modelica. In *International Modelica Conference*, pages 517–526, 2006.
- M. Wetter, W. Zuo, and T.S. Nouidui. Modeling of heat transfer in rooms in the Modelica "Buildings" library. In *Proceedings of Building Simulation 2011*, pages 1096–1103, 2011.
- M. Wetter, W. Zuo, T. Nouidui, and X. Pang. Modelica Buildings library. *Journal of Building Performance Simulation*, 7(4):253–270, 2014.
- M. Wetter, M. Bonvini, and T.S. Nouidui. Equation-based languages – a new paradigm for building energy modeling, simulation, and optimization. *Energy and Buildings*, 117:290–300, 2016.