Feedforward-Feedback Controller Using General Regression Neural Network (GRNN) for Laboratory HVAC System: Part III—Temperature Control—Heating

Osman Ahmed, Ph.D., P.E.
Member ASHRAE

John W. Mitchell, Ph.D., P.E.
Fellow ASHRAE

Sanford A. Klein, Ph.D.
Fellow ASHRAE

ABSTRACT

The simulated performance of the combined feedforward-feedback controller is compared to the performance of the feedback only (proportional-integral) approach for temperature control during heating. The heating sequence occurs due to a high room ventilation load caused by the fume hood openings. Local heating is provided by reheat coils. The simulated results show that the combined controller provides stable and accurate control under different operating conditions. The feedback controller only performs better than the combined controller at the operating condition for which it was tuned.

INTRODUCTION

This is the last part of the three-part paper focusing on a combined feedforward-feedback controller for variable-air-volume (VAV) laboratory HVAC systems. This part deals with the heating sequence commonly found in a VAV laboratory and compares the simulated performance of the current feedback approach to the proposed combined controller. Again, the combined approach outperforms the feedback controller considering wide operating conditions and different HVAC component characteristics (i.e., heating coil, damper, and valve).

This paper begins with a description of the heating sequence. Next, implementation of a combined approach for heating is discussed, followed by simulation and results. At the end of the paper, a general conclusion is drawn based on the observed results from all three parts of this study. Future research is also recommended.

In most VAV applications, the supply air is discharged into the laboratory space at a constant temperature of 55°F (12.77°C). Based on the normal design cooling load, the supply volumetric flow rate is selected to maintain the specified room temperature, usually a value between 70°F (21.11°C) and 75°F (23.89°C). In order to maintain the differential pressure, it is necessary that the minimum total lab exhaust exceed the supply flow rate due to the cooling demand. However, when the lab exhaust suddenly increases due to a fume hood sash opening, the supply flow rate also increases accordingly. The new supply flow rate at a constant 55°F (12.77°C) may exceed the requirement of the cooling demand. The room temperature, therefore, may drop below the set point. This sequence requires the local reheat valve to open and increase the supply air temperature to keep the room temperature set point. The coupling between room pressure and thermal constraints is complex.

For the purpose of this evaluation, two separate disturbance sequences are considered for heating. In the first sequence, the disturbance is caused by the sudden increase in the fume hood exhaust due to the sash opening from minimum to full open position. The laboratory exhaust is increased from an initial value of 450 cfm to a final exhaust of 2450 cfm (1156.40 L/s). The second control sequence considered for heating has a change in the space internal load. This sequence assumes that the lab initially has an initial maximum lab exhaust flow of 2450 cfm (1156.40 L/s) and a corresponding supply flow that maintains the room temperature and differential constraints. The total lab exhaust is then decreased to 1300 cfm (613.60 L/s) and, at the same time, an internal load of 13.20 Btu/ft^2·h (41.58 W/m^2) is generated. As a result, the space needs partial heating.

IMPLEMENTATION OF COMBINED CONTROLLER FOR HEATING

During heating, both the supply flow set point $\dot{v}_{a|sp}$ and the coil discharge air temperature set point, $T_{a,o|sp}$ need to be

Osman Ahmed is a senior principal engineer at Landis and Staefa, Inc., Buffalo Grove, Ill. **John W. Mitchell** and **Sanford A. Klein** are professors at the Solar Energy Laboratory at the University of Wisconsin, Madison.

determined. The set point is calculated based on the difference between the room temperature set point and actual value. The controller then generates a control signal C_s that is sent to the valve/actuator. Similar to the action of a damper, the valve modulates the water flow rate through the coil, $\dot{v}_f$, depending upon the valve authority, a, and the magnitude of the control signal, C_s. The inputs are the given air, $\dot{v}_a$, water flow rates, $\dot{v}_f$, the water inlet temperature, $T_{f,i}$, and supply air inlet temperature, $T_{a,i}$. The coil outputs are water and air outlet temperatures, $T_{f,o}$, and $T_{a,o}$, respectively. Often the coil output is expressed as dimensionless variable, R, which can be viewed as a coil effectiveness. R can be expressed as

$$R = \frac{T_{a,o} - T_{a,i}}{T_{f,i} - T_{a,i}}. \quad (1)$$

For heating, the physical characteristics are identified and inverted for use in control. The valve identification process is similar to the supply and general exhaust dampers described in Parts I and II (Ahmed et al. 1998a, 1998b). For the coil, the general regression neural network (GRNN) produces an output of water flow rate, $\dot{v}_f$, for the desired coil discharge air temperature set point, the $T_{a,o|sp}$, and given supply airflow rate set point, $\dot{v}_{a|sp}$. The GRNN captures the coil characteristics based on measured data. The feedback controller provides a residual control signal in order to maintain the desired coil discharge air temperature set point. The schematics of coil identification and control are shown in Figure 1. Again, the identification of the entire heating process is split into the coil and the valve in order to facilitate the implementation of the combined loop. The details of the identification and control of the heating process are discussed by Ahmed et al. (1997, 1996). The implementation scheme for the supply flow control loop remains the same as discussed in Parts I and II of this study. Both the supply flow and discharge air temperature set points are determined using the same physical models and method as described in Part II (Ahmed et al. 1998b) for temperature control—cooling sequence.

SIMULATION AND RESULTS

Temperature Control: Heating Sequence

For the heating sequence, the simulation sample time is chosen to be two seconds. Published literature (Underwood 1989; Gartner and Harrison 1965; Pearson et al. 1974; Nesler and Stoecker 1984) reports that heating coils respond in ten seconds to a temperature increase of 30°F (16.66°C)and in 40 seconds to a temperature increase of 5°F (2.78°C) across the coil. A coil response time of 20 seconds is needed to meet the rapid heating requirement in a lab environment.

The coil is modeled as a first-order system with dead time (Ahmed et al. 1997; Ahmed 1996). Therefore, the expression for the coil response time will be the same as for an actuator as shown in Part II for a damper actuator (Ahmed at al. 1998b, equation 13). For a given coil discharge air temperature set point, the response of discharge air temperature will vary with

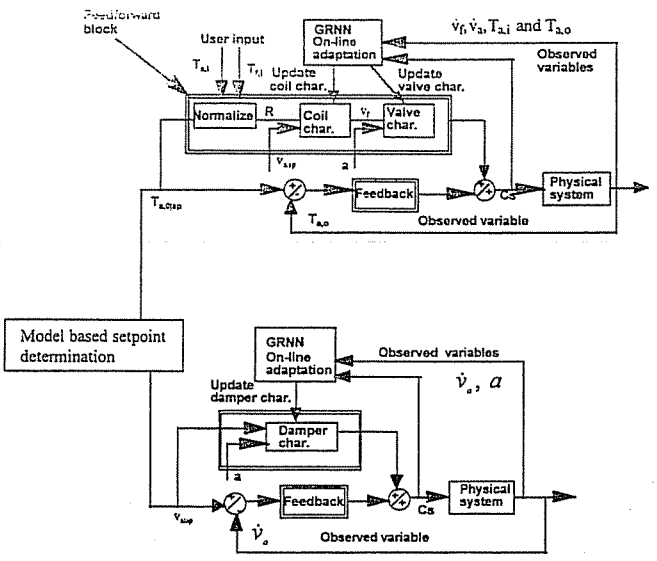

Figure 1 Implementation of combined controller for the heating coil.

the ratio of sample time to coil time constant. This ratio was studied, and a value of 0.40 was chosen as reasonable in order to obtain a desired coil response for reaching the discharge air temperature set point at full flow. The coil time constant, τ_{coil}, was determined from a simulated open loop coil response curve. The method of determining the coil time constant follows the same procedure as discussed in Part I (simulation section). The calculated value of coil time constant was found to be about ten seconds. A sample time of two seconds, therefore, means that about five samples are taken per coil time constant, which is adequate, as explained in Part I (simulation section). The choice of a smaller sample time compared to that for the cooling sequence (ten seconds) means that more samples are now available. However, the total time for simulation for the heating sequence is decreased to 20 minutes compared to 80 minutes for the cooling sequence. A range of 20 minutes is adequate considering that there is only one disturbance (i.e., increase in lab exhaust) for heating as opposed to the two disturbances for cooling. The tuning of the valves for the PI and combined feedforward and feedback (FFPI) controllers are done in the same way as described before with the pressure control sequence. The tuning parameters are shown in Table 1. The time derivative of pressure is also ignored in the simulation, as was the case in the cooling sequence.

In the heating sequence for which the ventilation load largely dictates the heating demand, the supply airflow discharge air temperature set point is determined using the steady-state mass, energy, and infiltration equations. The internal thermal load remains constant and does not need to be predicted. The supply flow rate set point required to satisfy the room differential pressure requirement is determined using the same set of steady-state equations. The model-based setpoint

TABLE 1
Table of Controller Gains

Control Sequence	Control Equipment	FFPI Controller		PI Controller	
		P_g (Control Signal/Error)	I_gS_t (Control Signal/Error)	P_g (Control Signal/Error)	I_gS_t (Control Signal/Error)
Heating	Supply Damper	3.0 *e*(-6)	4.9 *e*(-6)	3.0 *e*(-6)	2.5 *e*(-5)
	Coil Valve	0.00057	0.00019	1.04	0.00780

determination has been elaborately discussed for the temperature control cooling sequence in Part II (Ahmed et al. 1998b). In the case of cooling, the general exhaust and supply flow set points were calculated, whereas the supply flow and the discharge air temperature set points were calculated for the heating sequence. The control strategy for FFPI and PI approaches are shown in Figure 2. The control strategies are similar to that for the temperature control of the cooling sequence. In Figure 2, the model-based setpoint predictor determines the supply flow and discharge air temperature set point. The supply flow set point is fed into the feedforward (FF) loop for supply flow control while the coil loop receives the discharge air temperature setpoint information. The PI control loop for heating with a ventilation load only is similar to FFPI without the FF blocks.

The control strategy for the FFPI and PI controllers for the cooling sequence with the room load shown in Figure 2 is very similar to the previous sequence without the room load. However, the model-based predictor now includes the predicted load in order to determine supply flow and temperature set points. Instead of using the discharge air temperature set point as in the ventilation disturbance, the room temperature set point is used in the feedback block for both FFPI and PI controllers.

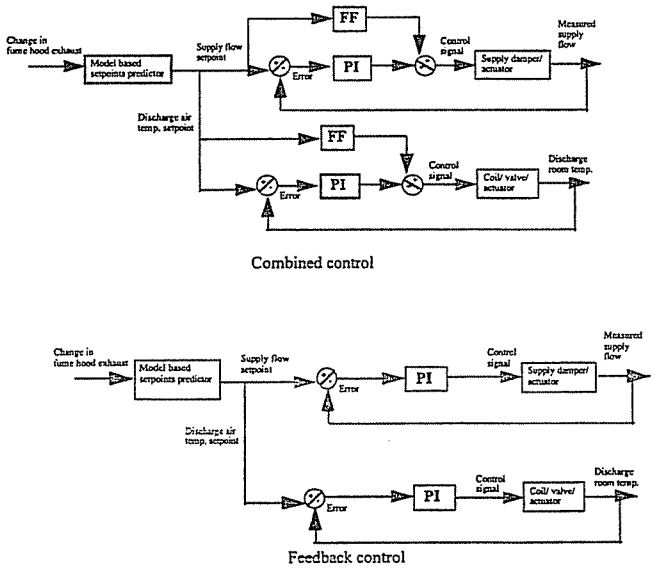

Figure 2 Schematics of controllers for temperature control—heating sequence due to ventilation.

If the predicted load is based on the discharge air temperature set point, any error associated with it propagates through the load calculation and makes the error even larger for the discharge air temperature in the next sample time. This produces an unstable response. This was observed repeatedly during simulation, and it was decided to use a more direct control variable (i.e., the room temperature and its constant set point.) in the control loop.

The PI controller block shown in Figure 2 appears to have two uncoupled loops: one for pressure and the other for temperature. However, in reality they are coupled since the current samples of room temperature and the discharge air temperature set points are used to determine the predicted load, which is then used to calculate the supply flow and temperature set points for next sample.

RESULTS AND DISCUSSION

Cases H1, H2, and H3: Ventilation Only

As in the case of cooling, three damper/valve characteristics are considered for this sequence. Identical characteristics are chosen for the damper and valve for each sequence. The first case, H1, considers a linear damper/valve with an authority of 1.0. Case H2 illustrates results for a linear damper/valve having an authority of 0.01, while the last case, H3, presents the results for a nonlinear damper/valve and an authority of 0.01. The system was initially at equilibrium at the disturbance suddenly applied at time equals to 20 seconds.

The plots for cases H1, H2, and H3 are shown in Figures 3, 4, and 5, respectively. The time response of the room and discharge are temperature variables that are usually measured; these are shown in the plots. In general, both PI and FFPI controllers work well for all control sequences. The good performance of the PI controller for all control sequences is expected because the operating condition requires the supply flow damper and the heating valve to open fully to meet the ventilation load caused by the increase in lab exhaust. The normalized control signals required to open dampers and valves are almost equal and close to unity irrespective of damper/valve characteristics. The room temperature falls below the set point initially but quickly recovers and reaches the set point. The simulation runs for about 20 minutes to ensure that the controller reaches the steady state without any oscillation or stability. The discharge air temperature quickly reaches the set point as determined by the mechanistic-based steady-state model, as explained in Part II (Ahmed at al.

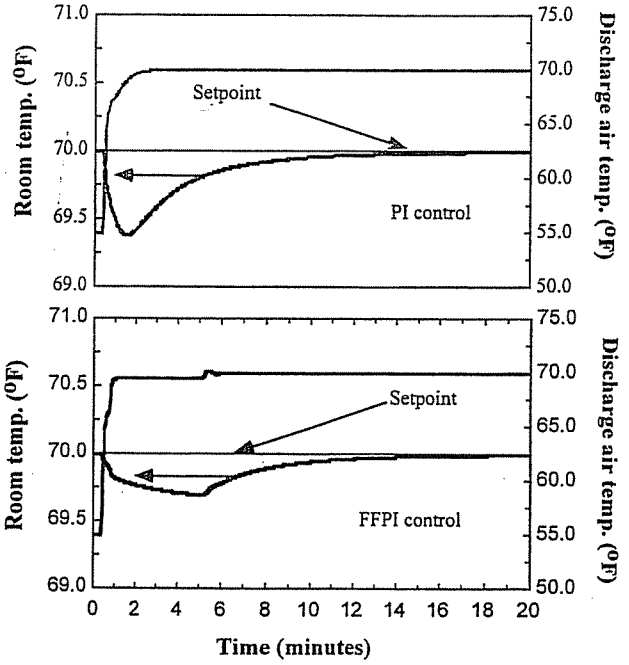

Figure 3 Dynamic response of room and discharge air temperature for control sequence H1.

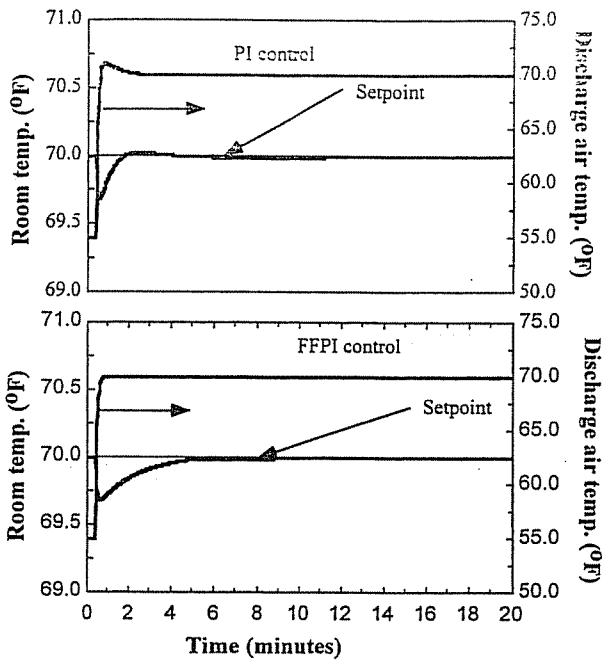

Figure 4 Dynamic response of room and discharge air temperature for control sequence H2.

1998b). The method of determining set points using physical models proved to be viable, as illustrated repeatedly in both pressure and temperature control sequences.

The initial undershoot in room temperature is expected as the sudden increase in room total exhaust causes the supply flow at 55°F (12.78°C) to increase. The increased supply flow rate is in excess of the required amount to offset the room thermal load. The room temperature thus falls before the coil water valve can open to provide heating and offset the sudden increase in room ventilation load. The initial undershoot in room temperature is less than 0.5°F (2.78°C) for both PI and FFPI controllers for cases H1 and H2. In the case of H1 for a linear damper/valve, the room temperature falls over 0.5°F (2.78°C) for the PI controller, whereas the undershoot is about 0.25°F (1.39°C) for the FFPI controller.

The recovery in room temperature in case H1 for both controllers is more gradual compared to case H2, as expected, since a quick opening damper/valve (i.e., authority 0.01) is used. In the case of an FFPI controller, the temperature continues to fall until the PI is activated (after about five minutes), which provides a residual control signal in addition to the FF signal to compensate only a small error of about 0.25°F (1.39°C) or less. This behavior demonstrates that the FF controller alone can provide a significant portion of the control signal that is required to achieve stability and accuracy.

In case H3, the PI controller produces a larger undershoot in room temperature compared to cases H1 and H2. The room temperature also shows an overshoot in the case of the FFPI controller. The unusual trend is explained with the help of Figure 6; the figure shows the response of temperatures for

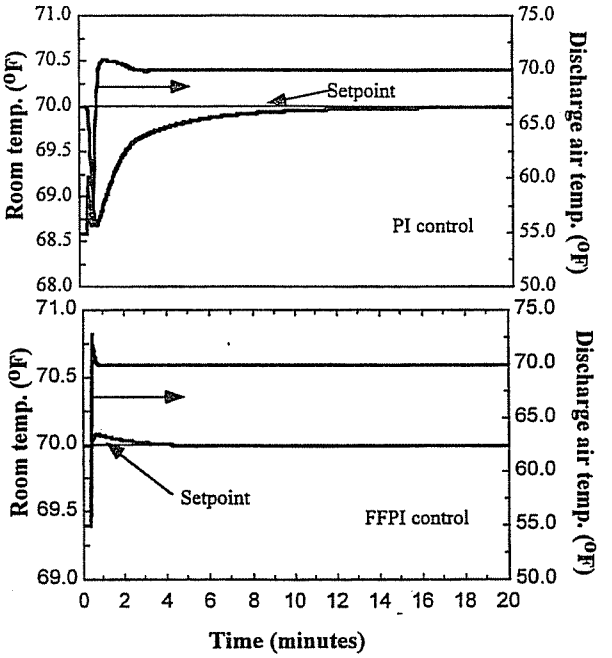

Figure 5 Dynamic response of room and discharge air temperature for control sequence H3.

only the first two minutes. With the PI controller, the discharge air temperature produces an initial oscillation and there is a slow response in achieving the desired set point. The room temperature response, therefore, takes longer and produces a larger undershoot before reaching the set point. The explanation for slow response is similar to what has been observed and

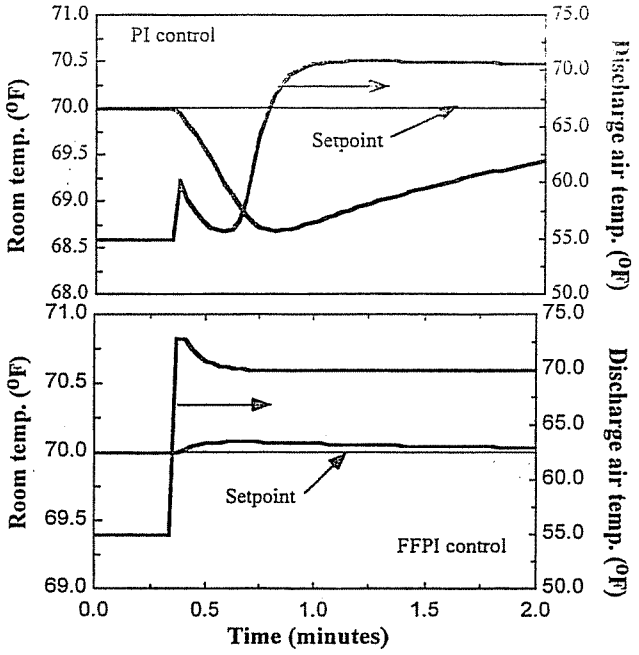

Figure 6 Initial response of room and discharge air temperature and predicted room load for control sequence H3.

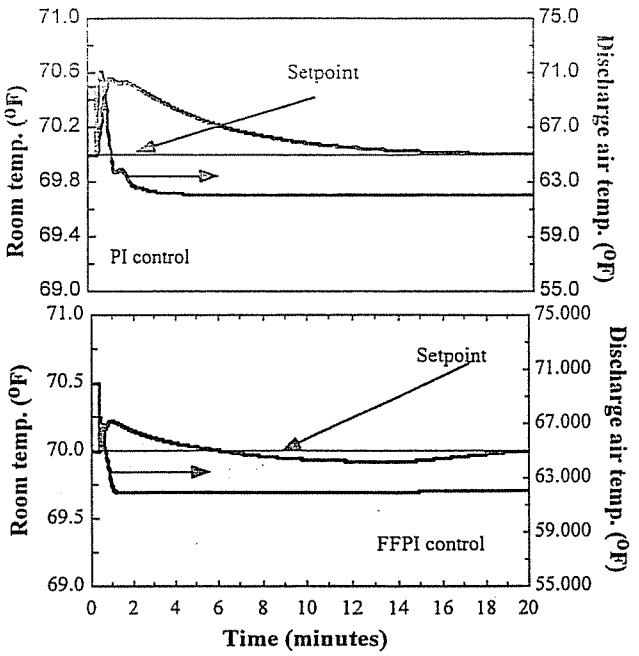

Figure 7 Dynamic response of room and discharge air temperature for control sequence H4.

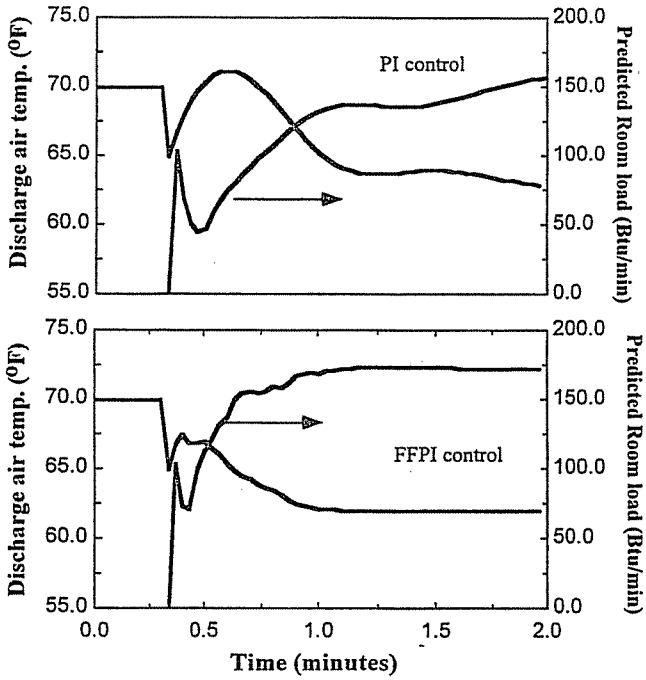

Figure 8 Initial response of discharge air temperature and predicted load for control sequence H4.

explained before in the case of pressure control sequence P3 (Part I, Ahmed et al. 1998a), which used an identical linear damper/ valve with an authority of 0.01. As a result, the damper/valve opens quickly, exhibiting its nonlinear characteristics.

The discharge air temperature overshoots considerably for an FFPI controller, which causes the room temperature to increase. The overshoot is caused again by the slow response of the nonlinear supply damper characteristics. Since the supply flow responds slowly, the room does not receive sufficient heating. As a result, the discharge air temperature set point continues to increase to offset the increasing heating load. Finally, when the supply flow reaches its set point, the room is overheated due to the higher discharge air set point.

Cases H4, H5, and H6: Ventilation with Heat Load

The last control sequence considers a sudden increase in heating load with the lab exhaust at mid-range. Both the heating valve and supply damper operating points are in the middle of their respective end-to-end stroke. The PI controller underperforms for nonlinear damper/valve characteristics and mid-operating points. On the other hand, the FFPI controller provides stable and accurate control for all three cases.

Cases H4, H5, and H6 consider a linear valve, a linear valve with an authority of 0.01, and a nonlinear valve (W_f = 0.5) with an authority of 0.01, respectively. The corresponding plots are shown in Figures 7, 9, and 10. Figure 8 shows the response for a short-time interval for a linear damper/valve. The general trend is similar in all cases. As the sudden load is imposed, the room temperature increases before the physical model can determine that less heating is needed and signal the control system to close the valve. The discharge air temperature then settles to a new steady-state value.

Figure 7 shows that both PI and FFPI controllers work very well for the linear damper/valve. The response of the

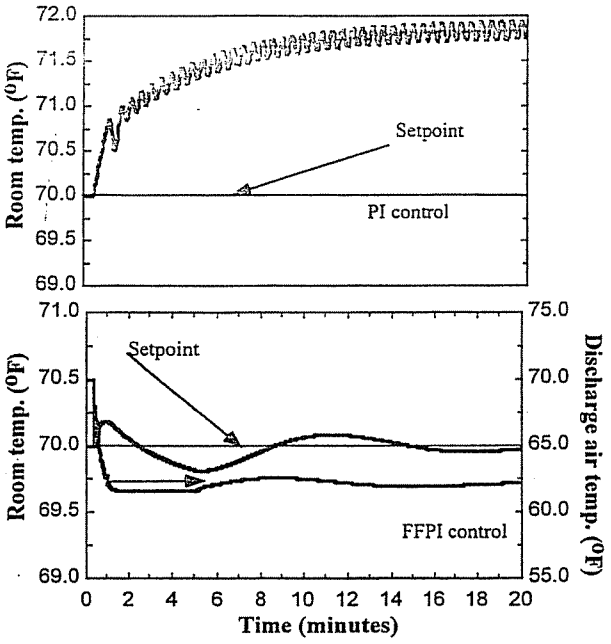

Figure 9 Dynamic response of room and discharge air temperature for control sequence H5.

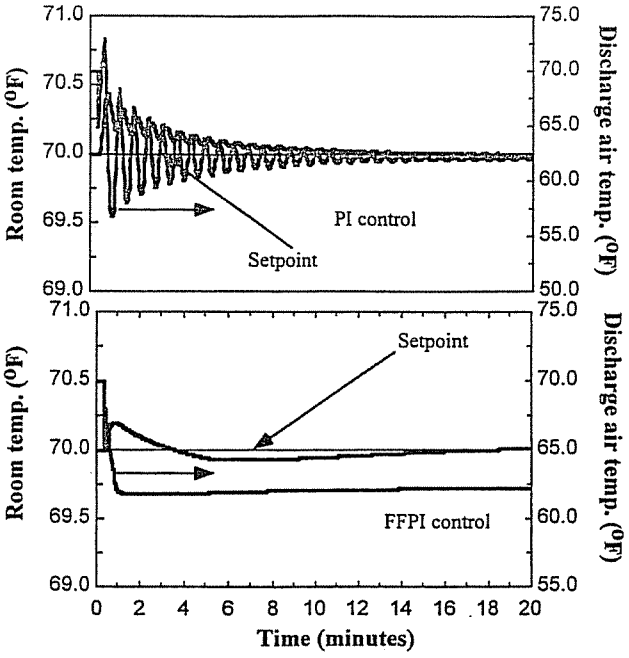

Figure 10 Dynamic response of room and discharge air temperature for control sequence H6.

discharge air temperature and predicted load shown in Figure 8 for an initial time period of two minutes reveals some unexpected transient behavior. The predicted load for both PI and FFPI controllers initially increases very rapidly followed by a dip and then a gradual increase to reach the steady-state value. The steady-state value of the predicted load is 167 Btu/min and 169 Btu/min (0.815 W/min and 0.825 W/min) for PI and FFPI controllers, respectively, compared to the actual 165 Btu/min (0.805 W/min) of generated load. The small difference between predicted and generated load is due to the heat transfer from the wall surface that is maintained at a slightly higher temperature than the room.

The initial sharp increase in the predicted load from zero is caused by the steady-state load using values that are determined of different variables at the sample time preceding time = 0. A sudden increase in heating load causes the discharge air temperature set point to jump, which results in a decrease in load. As the initial disturbance becomes steady, the predicted load then continues to rise steadily to reach the steady state.

The plots in Figures 9 and 10 demonstrate the deficiency, once again, of the PI controller as the operating points shift away from the tuning point. The PI controller produces poor control, and for case H5, continues to operate the control valve in an oscillatory fashion. The room temperature then oscillates. In case H6, the oscillation amplitude decreases from a large initial amplitude but is still present after 20 minutes. A similar yet reverse trend is noted when the response in cooling sequence C5 (Part II, Ahmed et al. 1998b) is compared to case H5 and the response of cooling sequence C6 (Part II, Ahmed et al. 1998b) is compared with the case H6. The similarity is expected as cases C5 and H5 use a linear damper with an authority of 0.01 while C6 and H6 use a nonlinear damper (W_f = 0.5) with an authority of 0.01. All of the cases use the predicted thermal load in the model-based setpoint predictor, but the load suddenly increases for C5 and C6 while the load decreases in the heating sequence. The FFPI controller for both cases H5 and H6 again show stable and accurate control.

SUMMARY AND RECOMMENDATIONS

The cases for pressure and temperature control sequences are selected with an objective of evaluating controller performance over a wide range of operating conditions and different equipment characteristics. The FFPI controller performs well for each case and demonstrates stable and accurate control. In contrast, the PI controller fails to perform adequately except at the operating condition that was used in tuning. The FFPI is able to respond quickly to any setpoint change due to the presence of the feedforward element, while the feedback loop provides stability and eliminates residual error between the set point and simulated values. As the damper characteristics shift from the assumed initial linear damper due to change in authority, the advantage of FFPI over the feedback loop becomes apparent. The FFPI controls rapidly and accurately over a wide range of operation even with significant change in damper characteristics. The FFPI used single tuning parameters for its feedback portion, and, thus, control is robust.

In an FFPI controller, a significant portion of the control signal is derived from the FF part while only a very small signal is needed from the PI part. The PI serves to improve the controller accuracy in reaching the desired set point and eliminates offset. As a result, the PI portion of the FFPI control loop only handles a small amount of error irrespective of

change in the system characteristics and operating conditions. This makes tuning of the FFPI very simple, and retuning is not necessary, which is often the case with the PI control loops.

The FF part requires only a single smoothing parameter to be estimated, which can be held constant for most of the HVAC processes (i.e., dampers and valves) because the amount of data needed to identify the characteristics of such processes do not vary significantly. Finally, the FF portion of the combined approach can adapt to changes in system characteristics, which is a feature that provides the correct control signal and diminishes the role of the feedback component.

A comparison between the proposed FFPI and a state-of-the-art PID control system for VAV was used to highlight the advantage of FFPI from the commissioning and operation point of view. Figure 11 shows the schematic of a total integrated control system for a VAV lab HVAC system for both PID feedback and proposed FFPI approaches. The feedback approach has a total of five coupled PID loops with a total number of 15 gains that have to be field tuned and adjusted. Even if all PID controllers are considered to be PI only, there are still ten tuning parameters. These coupled loops are difficult to tune and are usually tuned at a fixed operating condition. Retuning may be necessary if the operating condition shifts considerably, which is common in a lab environment. In general, the poor performance of the coupled feedback loops proves costly for the lab owner.

In contrast, by shifting the burden of performance from the PI to the FF component, the tuning may become a non-issue for the proposed FFPI controller. The ability of the FFPI to work over a wide range of operating conditions and component characteristics is very promising.

There is to no need for retuning to produce stable, robust, and accurate control. The FFPI implementation in a real

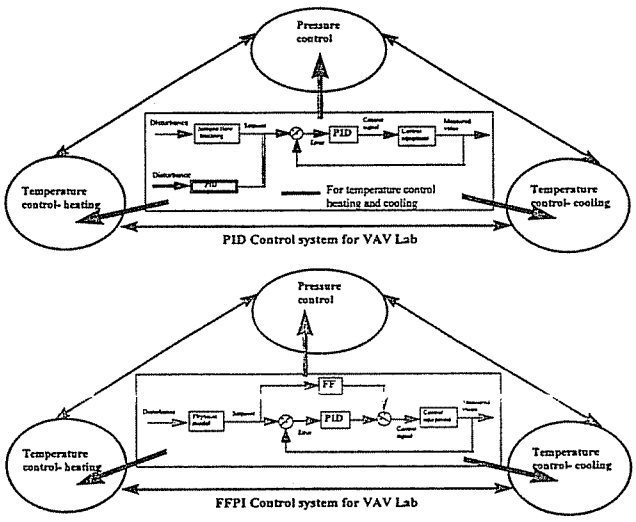

Figure 11 *Schematic of a control system for a variable-air-volume laboratory.*

controller will translate into simple commissioning and better performance and, as a result, add value to the customer.

Furthermore, the FFPI controller does not need any additional sensors, added memory, or processing capability. The sensors that are normally used in a current control system will be able to provide the required information for HVAC control equipment (i.e., valve and damper identification). The use of the general regression neural network (GRNN) and model-based setpoint predictor are simple algorithms and do not need significant controller memory. Therefore, the FFPI may be implemented without changing or adding any hardware/sensor configuration. The only investment will be needed for algorithm development. The simplicity associated with the FFPI will be a real cost savings to the customer.

Based on the results of this study, the following recommendations can be made.

1. The present study completes a "proof of concept" phase. Verification of the actual implementation of combined feedforward and feedback loops in a laboratory VAV HVAC system will be the next step. The combined approach may be implemented using existing state-of-the-art hardware/software found in a building control system. The performance then needs to be compared with the current feedback approach in practice. Besides the control performance, the tuning, troubleshooting, maintenance, ease of commissioning, and operation will be key success factors.
2. The pressure and temperature control sequences used to test the controllers were considered separately in this research. In reality, however, the sequences will be coupled and overlapping. For example, while the increase in fume hood exhaust flow may require heating, the increase in the rate of internal heat generation may dictate more cooling. It will be necessary to test the controller with simultaneous cooling and heating demands.
3. In addition to the space temperature, space humidity might be included as an additional control variable. The humidity control is critical in the process and manufacturing labs but is seldom being controlled in a common research lab.
4. A supervisory controller is needed that will command and coordinate the actions of the supply flow, general exhaust flow, and supply air temperature control loops. The focus of the supervisory controller will be to manage these different control loops in order to eliminate contradictions between them and to ensure a smooth transition from one control sequence to another.
5. In the simulations, the control signal from the feedback approach is added to the feedforward control signal after a certain time delay following the initial disturbance. The delay time was selected for each control sequence during tuning of the FFPI controller and was kept constant for all cases of a specific control sequence. A scheme may be developed that combines the feedforward with feedback approach automatically. The method of combination may

be based on the delay time, as was the case in this study, or on the achievement of steady state in terms of a process variable (e.g., flows, room temperature, pressure, etc.). The combination method preferably should be a part of the commissioning process.

6. The success of any identification scheme, including GRNN, depends on gathering quality data from a real system. A data acquisition system with post-processing capability is required for this purpose. In addition, a method needs to be adapted to update the current data stored and used by the GRNN. The adaptation is absolutely necessary in order to capture the effect of change in operating conditions on the characteristics. The ability to adapt is critical to reducing the retuning effort, which in turn will require minimum troubleshooting and will increase operational efficiency.

7. Any hysteresis associated with damper and valve actuator action has been assumed negligible. Damper friction and actuator sticking are also neglected. A modification of the identification block may be suggested to include hysteresis as well as other damper/actuator properties with the HVAC control equipment. In order to include hysteresis, one approach may be to develop two separate GRNNs for each piece of control equipment. One would identify the equipment characteristics when the control signal increases and the other one when the control signal decreases. A simple algorithm may select one GRNN over the other by noting the change in the direction of a setpoint change from the current value.

8. The initial conditions for the simulation of the different control sequence were for steady and constant operation. The boundary conditions are the supply and general exhaust duct pressures, fume hood exhaust flow rate and pressure, and temperature of the adjacent space. In reality, however, these conditions will have noisy data and fluctuate. Further simulation may be recommended to accommodate such noise and prove the control system's ability to handle noisy data.

9. Reliability of the identification block in the feedforward control has to be investigated. The risk of poor identification and its effect on control need to be assessed. Clustering and filtering of data for identification may be needed in order to increase accuracy of the feedforward control signal.

REFERENCES

Ahmed, O., J.W. Mitchell., and S.A. Klein. 1998a. Feedforward-feedback controller using general regression neural network (GRNN) for laboratory HVAC system: Part I—Pressure control. *ASHRAE Transactions* 104(2).

Ahmed, O., J.W. Mitchell., and S.A. Klein. 1998b. Feedforward-feedback controller using general regression neural network (GRNN) for laboratory HVAC system: Part II—Temperature control—Cooling. *ASHRAE Transactions* 104(2).

Ahmed, O., J.W. Mitchell., and S.A. Klein. 1997. Feedforward-feedback HVAC controller using general regression neural network. IFAC Symposium on Artificial Intelligence in Real Time Control, Kuala Lumpur, Malaysia, September.

Ahmed, O. 1996. Model-based control of laboratory HVAC system. Ph.D. thesis, University of Wisconsin-Madison.

Ahmed, O., J.W. Mitchell, and S.A. Klein. 1996. Application of general regression neural network (GRNN) in HVAC process identification and control. *ASHRAE Transactions* 102(1).

Gartner, J.R., and J.L. Harrison. 1965. Dynamic characteristics of water-to-air crossflow heat exchangers. *ASHRAE Transactions* 71(1): 212-224.

Nesler, C.G., and W.F. Stoecker. 1984. Selecting the proportional and integral constants in the direct digital control of discharge air temperature. *ASHRAE Transactions* 90(2b).

Pearson, J.T., R.G. Leonard., and R.D. McCutchan. 1974. Gain and time constant for finned serpentine crossflow heat exchangers. *ASHRAE Transactions* 80(2).

Underwood., D.M. 1989. Response of self-tuning single-loop digital controllers to a computer-simulated heating coil. *ASHRAE Transactions* 95(2).

NOMENCLATURE

a	=	damper/valve installed authority
FF	=	feedforward
FFPI	=	combined feedforward and feedback
GRNN	=	general regression neural network
HVAC	=	heating, ventilating, and air conditioning
I_g	=	integral gain constant in PID controller
PI	=	proportional-integral
PID	=	proportional-integral-derivative
P_g	=	proportional gain constant in PID controller
r	=	normalized actuator position (0-1)
r_a	=	command actuator position
S_t	=	sample time, seconds
t	=	time
T	=	temperature, °F (°C)
VAV	=	variable air volume
$\dot{v}$	=	volumetric flow rate, ft^3/min (m^3/min)
W_f	=	nonlinear valve/damper parameter
w.c.	=	inches of water column gauge
w.g.	=	inches of water column gauge

Greek Symbols

τ	=	time constant, seconds

Subscripts

a	= air
ad	= adjacent space
coil	= coil
e	= exhaust
ex	= general exhaust
f	= fluid (water)
fh	= fume hood
s	= supply
sp	= set point
ss	= steady state
st	= thermostat