Optimized Response to Thermal Sensation Complaints in Buildings

Rodney A. Martin*

Department of Mechanical Engineering
University of California, Berkeley

A report submitted in partial satisfaction of the

requirements for the degree of

Master of Science

in

Mechanical Engineering

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Research Advisor:

Professor David M. Auslander

Co-Advisor:

Dr. Clifford Federspiel

*Supported in part by the Center for Environmental Design Research and the College of Engineering Graduate Academic Diversity Engineering program

Thanks to Carla M. Trujillo, Cliff Federspiel, & Prof. Auslander for their assistance and support

Contents

List of Figures

List of Equations

List of Tables

1 Abstract

The evolution of commercial building energy management systems that control HVAC systems within the past decade has been focused on better energy management and reduced operating costs. Due to technological changes in computing, including processors, databases, network communications, and control systems applications, much greater capability has been employed to operate “smart” buildings. However, despite these vast advances in technology, fundamental control issues still remain; occupants still complain about comfort conditions. It has been found that current industry methods are adequate at best and random at worst (Smothers, Haley, Fisher et al., 1999). Facility managers may be making more work for themselves by responding to uncomfortable and unsatisfied customers using ad-hoc approaches, while losing money due to the extra cost of responding to complaints inefficiently.

Methods to optimize the manner in which facility operators currently respond to building occupant complaints has not been researched until recently (Federspiel, 1998). Finding an alternative thermostat setpoint strategy using building occupant complaint trends and simple optimization methods is attainable. In fact, the results imply that modification of current strategies is paramount to reducing overall cost & improving customer satisfaction. Additionally, labor charges by facility operators responding to complaints can be reduced effectively, by integrating these optimal methods into current state of the art direct digital control systems, allowing for control system automation of setpoint changes. Federspiel suggests that by using these new strategies to minimize complaints, there is a $2-3 billion maintenance cost avoidance potential (“Statistical Analysis”, 1998, pp. 921-922). Therefore, ad-hoc or non-use of building occupant complaint information can be costly in the long run.

2 Introduction

2.1 Motivation

75% of all environmental complaints recorded in buildings are of the thermal sensation type, as opposed to humidity, air circulation, etc. Furthermore, 40% of these thermal sensation complaints occur when there are no faults in the servicing HVAC system. (Smothers, Haley, Fisher et al., 1999). Therefore, for simplicity this project will be based upon hot and cold complaints alone, when the building's HVAC systems are operating normally. The no-fault thermal sensation complaint handling problem can be posed as a feedback control problem, with building occupants acting as the sensors, and the thermostat setpoint strategy providing the control law, depicted in following schematic:

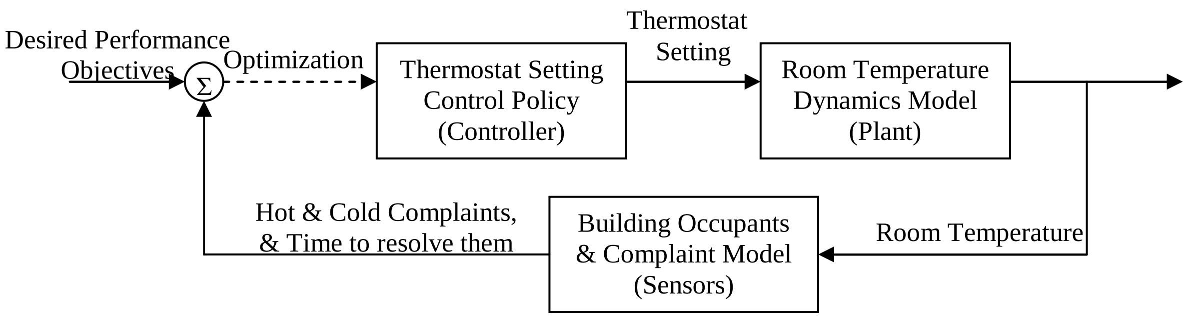

Figure 1- Building Occupants in Feedback Control Loop

The thermostat setpoint strategy shown above will be designed to replace current strategies by reaching desired performance objectives via optimization. One objective is to realize cost savings, potentially as much as $2 billion for the entire US stock of buildings, and $60 million for GSA buildings (Federspiel, "Statistical Analysis", 1998, pp. 921-922). An additional performance objective is to improve building occupant satisfaction by reducing response time to thermal sensation complaints.

2.2 Objective

The standard industry corrective action taken in response to thermal sensation complaints is to adjust the thermostat setting appropriately, leaving it there indefinitely (Smothers, Haley, Fisher et al., 1999). An alternative thermostat setting policy is to make an adjustment when a complaint is received, changing it back to the original setting after a certain period of time. This "react & setback" policy is parametrized by two variables: the magnitude of the setpoint change, and duration of time at the new setting. The goal is to find an optimal pair of values for magnitude & duration. These values can be obtained by running repeated simulations over a finite range of thermostat setpoint magnitudes & durations. As with any process, there must be a means of performance measurement. The performance measures: customer needs and cost, are the basis of this optimization. Customer needs are based on the responsiveness of facility operators in resolving a complaint. Because the investigation of complaints often leads to faults that need corrective action within the system, the period in which a building occupant is waiting for resolution is referred to as the "Fault Recovery Period". This project is investigating no-fault thermal sensation complaints only, so an alternative nomenclature is "Complaint Recovery Period" in lieu of Fault Recovery Period. These terms can be used interchangeably, in addition to their acronyms FRP and CRP. However for consistency, Complaint Recovery Period or CRP will be used most often to describe this metric. The cost measurement is based upon the number of complaints generated within a specific timeframe. Cost relates to complaints due to the general monetary resources allocated toward them, in time charged by facility operators, and general & administrative overhead costs associated with processing the complaint. To make an adequate comparison between current practice and optimized thermostat setting policies, the method of measurement needs to be unbiased. As such, the metrics described will need to be normalized. Customers' needs will be measured as the **average** complaint recovery period, or time per service call to resolve the condition causing the complaint. Cost will be measured as the **annual** number of hot & cold complaints recorded. For the optimized strategy, finding the magnitude & duration that minimizes both of these metrics is the primary objective. The subsequent analysis step is to compare the results found from the optimized strategy to the current practice metrics. All performance metrics come from computer simulation using C code, and the offline optimization/analysis steps are performed with MATLAB®. For the new strategy this results in a "static" thermostat setting policy, as opposed to a dynamic one that is adaptive and changes the parameters during simulation runtime. Figure 1 shows a dotted line under the word "Optimization" to illustrate the offline nature of control in the feedback loop. The following figure illustrates the similarities & differences between the two thermostat setpoint strategies:

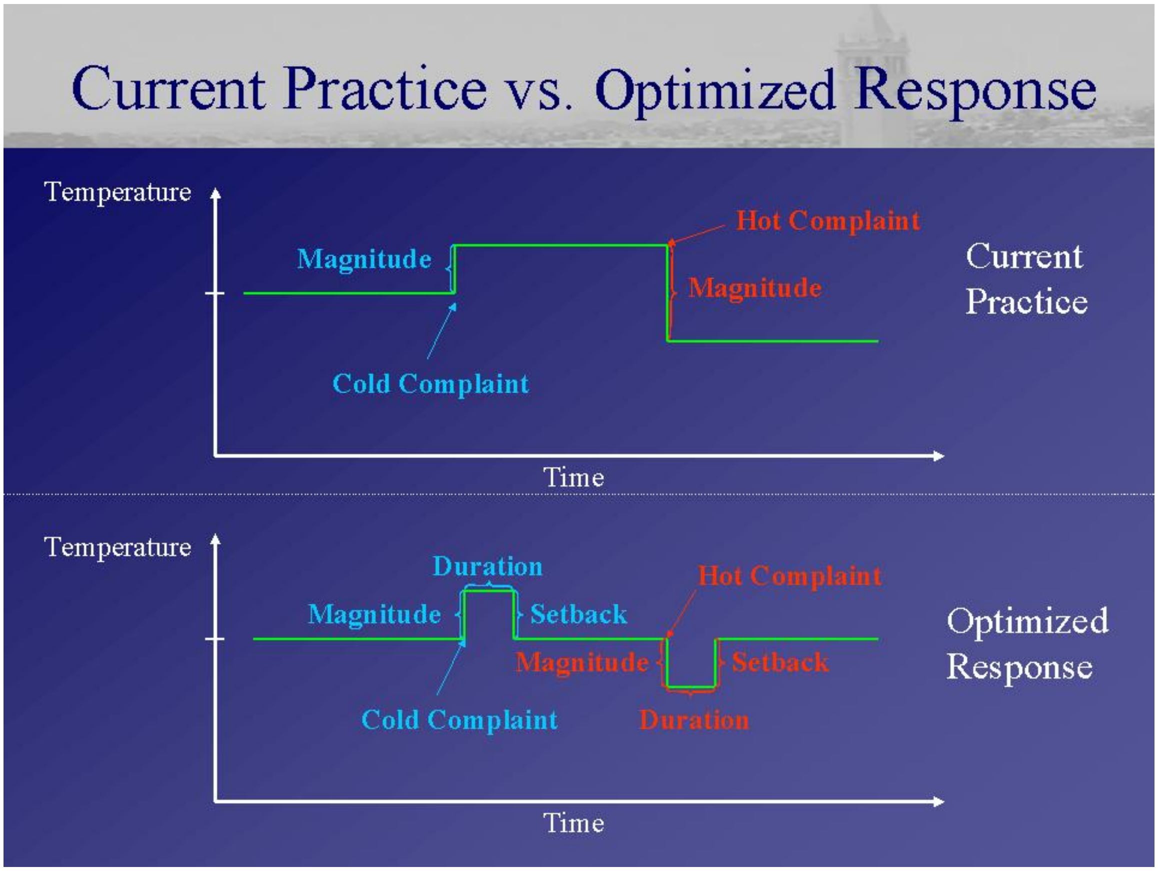

Figure 2 - Comparison of Thermostat Setpoint Control Policies

As seen in the figure above, the current practice strategy exhibits "typical" directional thermostat setting characteristics. In other words, the direction of the change is always as expected: if a hot complaint is received, the thermostat setting is decreased, and if a cold complaint is received the thermostat setting is increased. However, there is no typical thermostat setting behavior for the magnitude of the change. It is unknown ahead of time, and therefore modeled as a random process. Research indicates that industry strategies exhibit behavior correlating current setpoint changes to prior setpoint changes (Smothers, 1999). The trend reveals that if a prior setting moves away from some nominal setting, any subsequent change returns close to that setting. Additionally, setpoint changes tend to remain small in magnitude, with larger setpoint changes being less frequent. As a result, this information can be used to choose the appropriate probability distribution representing the statistical model of the industry response to complaints. This will be discussed in more detail later.

For the optimized strategy, duration is considered as well as magnitude, hence making it a two parameter thermostat setting control policy. The policy has two steps: the first step involves reacting to the thermal sensation complaint using the same direction as in the industry model. However the change occurs with a **specified** magnitude that is completely deterministic, as opposed to a statistical one. The second step is to adjust the thermostat setting back to its nominal value after a specified duration of time. The nominal setting has been optimized to yield the minimum annual number of complaints as well as the minimum average complaint recovery period. State logic machines deal with extenuating circumstances such as multiple complaints prior to setback, and simultaneous complaints of opposite types, etc. These circumstances will be discussed in detail in an ensuing section.

To summarize the similarities of the two strategies, it appears that the direction of the setpoint change should be opposite to the type of thermal sensation complaint received for both. Current practice policy differs from the new strategy because new settings have varying statistical magnitudes, as opposed to the same deterministic magnitude found to be optimal. The current practice thermostat setting only returns close to the nominal setting when a subsequent complaint is received, whereas the new strategy always returns to the nominal setting after an optimized duration of time regardless if there is a complaint or not. Finally, the duration associated with the new strategy is a persistently finite, optimized value, based upon the results of the offline optimization, whereas the industry strategy setting is changed for an indefinite amount of time until a complaint of the opposite type occurs. The main goals of using the optimized strategy are to reduce costs by minimizing the number of complaints, and speed up the time in which the complaint is resolved. As a result, occupant satisfaction should be improved by applying this new type of complaint handling. The optimized policy can be implemented via direct digital controllers, but also manually by facility operators.

2.3 Basic Elements of Building Operations

2.3.1 General Discussion

Simulation of complaints lies at the root of the optimization problem. Pragmatically, complaint data comes from building occupants as part of an archive, stored in a maintenance management system. Because building occupants are mobile, ubiquitous, intelligent, and possess many sensory modalities, they can act as excellent sensors in a control loop and provide multi-modal offline feedback measurements to quantify & qualify the current comfort level. Information derived from the comfort level of the occupants, traditional room temperature measurements, and archived maintenance information found in databases is useful in the development of new thermostat setting strategies such as the offline optimization method described earlier. These strategies can be used to accomplish better control over cost & customer satisfaction. The comfort level of building occupants is often measured via unsolicited complaints, as opposed to the use of a one-shot survey. Therefore, although technically discrete and stochastic, unsolicited complaints provide for a pseudo-continuous stream of building occupant feedback. The solicited survey-based method is a discrete, deterministic process controlled externally. Since complaints from building occupants are almost always expected, using the unsolicited method is an excellent way to capture “free” data. In this way, no provisions for alternative forms of measurement that might be more expensive and time-consuming are needed. However, because the receipt of unsolicited complaints is a random process, the frequency and measurement of complaints will directly affect the stability of the feedback control loop illustrated in Fig. 1.

This project can be extended beyond the implementation of a new thermostat setting control policy. It may include the integration of several different key information sources & interactive components within a building. As shown in Fig. 3, a collateral objective presents itself by using occupants in the feedback loop.

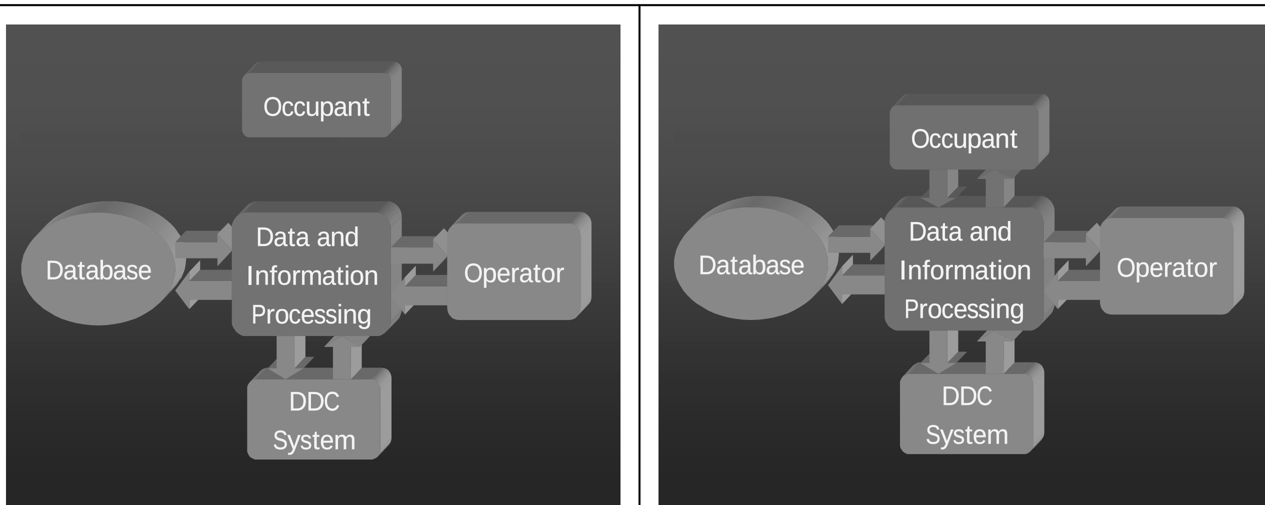

Before involving building occupants **No Direct System Integration**	*After involving building occupants* **System Integration**

Figure 3 - Occupant Feedback Inclusion

For simplicity, building operations has been broken into five key elements, as shown above. The data and information processing center is the core of the system, representing the computer network acting as the main conduit for moving information back & forth. The remaining four elements can communicate with the data & information processing center, providing for enhanced building system inter-operability. However, prior to the integration of building occupants, building operations are limited in functionality. As shown on the left, information entered by operators who log complaint data and maintain maintenance logs populate the database, hence their actions directly affect the DDC system as well. The DDC system also plays an interactive role, acting on external sensor readings and data stored in the central computer. Although this functionality seems adequate, it does not involve the building occupants, or capture the important benefits which result. As seen on the right, after the occupants are involved directly in the process, they also send information to the main computer used by the DDC system, in addition to information from the other elements. However, the relationship between building occupants and the data & information processing center is bi-directional. In light of this, an occupant interface comprised of front-end software is necessary for seamless operation, making the integration of all elements shown in Fig. 3 complete. Database driven maintenance management systems successfully being used in facility management today would satisfy the requirements for an occupant interface. These systems would be used to generate work requests due to incoming complaints, and to obtain information from occupants. They would also retrieve local temperature readings from field panels, sending information about the current status of the system back to the occupants.

2.3.2 Specific Application to Current Research

It is evident that this project has larger implications for the facility management domain. For this project in particular, only building occupant complaint temperatures are being investigated. Intrinsically, it is a reasonable assumption to make that generation of complaints & corresponding temperature recordings occur at the same time. In practice, this is not necessarily always the case. Therefore this study may eventually be extended to include parallel sensor readings. This would include temperature readings from operators at field panels, historical information from the database pertaining to corrective & preventative maintenance, and comfort level measurements from the building occupants. As a result the complete integration of all five elements would be realized, in which an expert system could be used to decide on the best control strategy. However, the main focus of this paper is to explore the more technical aspects of adding the new building occupant element into the system, and its potential interactions with the DDC system. Therefore, a modified diagram to describe this report's contents is as follows:

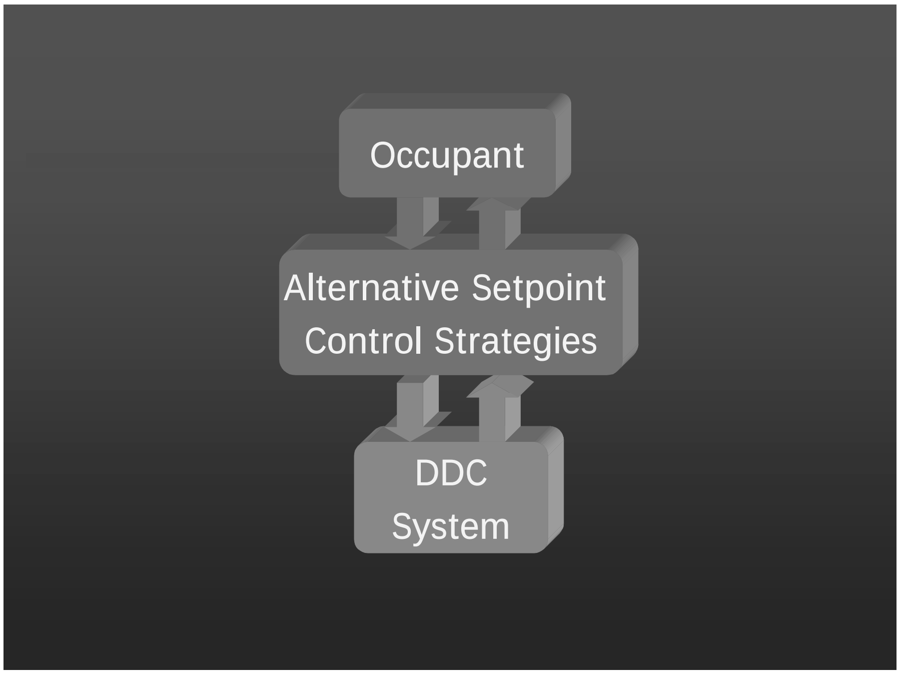

Figure 4 - Thesis Objective for Occupant Feedback Inclusion

For the aggregate control system shown in Fig. 1, performance metrics implicitly determine the cost of keeping the occupants in the space comfortable, and the response time required to meet the demands of their optimal comfort level. Although the details will be discussed later, a brief prelude is in order to allow for a clear understanding of how these metrics are measured in simulation. A simulated complaint is generated by a level crossing of two processes: one that determines the building temperature, the other determining the temperature at which a building occupant will complain. This might be best shown with a picture:

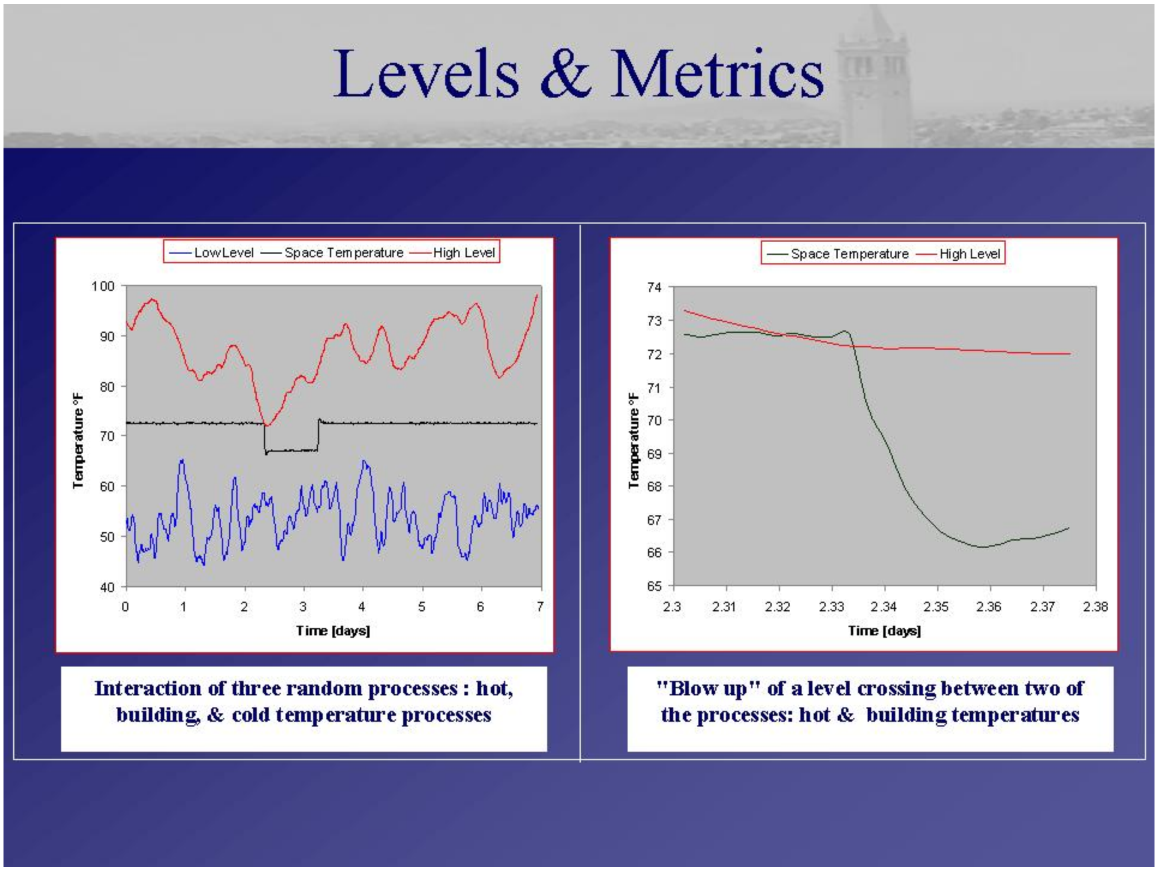

Figure 5 - Statistical Model of Complaint Behavior

As seen in Fig. 5, the picture on the left displays three distinct random processes. The top (red) line represents the temperature at which a building occupant will complain if they feel hot. This is similar to a level at which an alarm would ring if it were exceeded. The center (black) line represents the building temperature process, and bottom (blue) line represents the temperature at which a building occupant will complain if they feel cold. Again, this is similar to a level at which an alarm would ring if the building temperature falls below it. Obviously, all three processes are random in nature, using measured data and complex modeling techniques as the basis for their statistical parameters & outputs. The modeling & simulation of these processes will be discussed in more detail later. The second picture on the right is a blow-up of a crossing that occurs between the hot complaint process, and the building temperature process. The metrics discussed previously used for optimization in the thermostat setting control strategy are generated from graphical crossings such as these. On the right side of Fig. 5, the first intersection of the top (red) line and the bottom (black) line is a level upcrossing between the hot complaint process and the building temperature process, which in simulation represents the generation of a hot complaint. The second downcrossing has no significance in the real world, although theoretically it can be thought of as the end of the hot complaint recovery period (CRP). A cold complaint is not shown in Fig. 5, but can be similarly generated if a downcrossing of the building temperature process with the cold complaint process occurred. The end of the cold complaint recovery period would take place at the subsequent upcrossing between the cold complaint & building temperature processes.

Feedback of building occupant comfort information is triggered only when complaints arise, and so therefore unlike most control systems, poses a unique problem in measurement. The feedback of these measurements is pseudo-continuous, as mentioned earlier, but by nature exhibits discrete and stochastic characteristics. Additionally, the complaints registered are not always consistent. In general, measuring the comfort level of the building occupants is a challenge within itself. However, assuming that the complaints registered are an adequate indicator of building occupant discomfort, the information is still usable. The feedback of this information would not be geared directly towards changing the input of the plant itself, but rather indirectly via altering the setpoint. This forms a cascade control system where the comfort information is used to modify the heating and cooling setpoints of the terminal unit controls for the associated occupant space. The control action would be purely automated, the setpoint changed via a mechanism within the system.

Theoretically there are two loops in this control system. One loop controls the HVAC system component via feedback from the thermostat (i.e. temperature sensor, thermostat, and air handling unit that allows for appropriate cooling or heating of the spaces). As long as the setpoint to this system remains constant, then the response characteristics of this control system are relatively commonplace. However, with the addition of second “occupants loop” the problem becomes more challenging. Here the occupants are acting as intelligent, but not necessarily predictable sensors. Because the range of responses from a “human sensor” can vary from the typical to the pathological, the convergence of this control algorithm to its desired objective becomes an interesting problem for study. However, there is still only one system to be controlled, which is the HVAC system associated with the thermostat. A block diagram of this inner loop process is shown in Fig. 6:

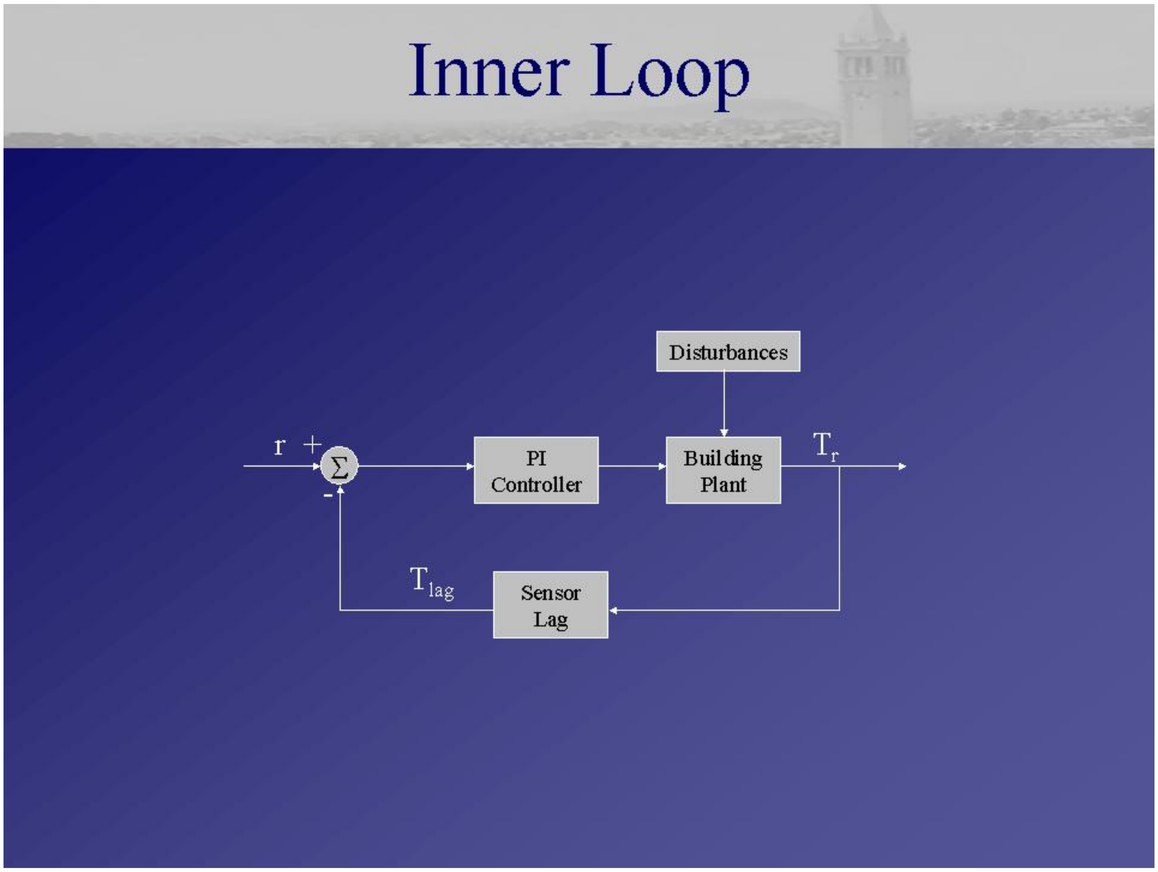

Figure 6 - Inner Loop Controller

In industry, it’s often found that there are time lags associated with the thermostats. Furthermore, PI controllers, or controllers with only two gains (two-term controllers) are almost always used in conjunction with them. On a macro level, there are additional “human sensors” measuring the comfort level in addition to temperature sensor feedback that is used in a simple PI control algorithm. They are located in the control loop back around to the desired setpoint. In this loop a setpoint control algorithm is used to alter the setpoint in such a way to allow for implementation of a thermostat setpoint control strategy exhibiting optimized performance metrics. A depiction of what the system might look like in block diagram format is a more detailed version of Fig. 1, as follows:

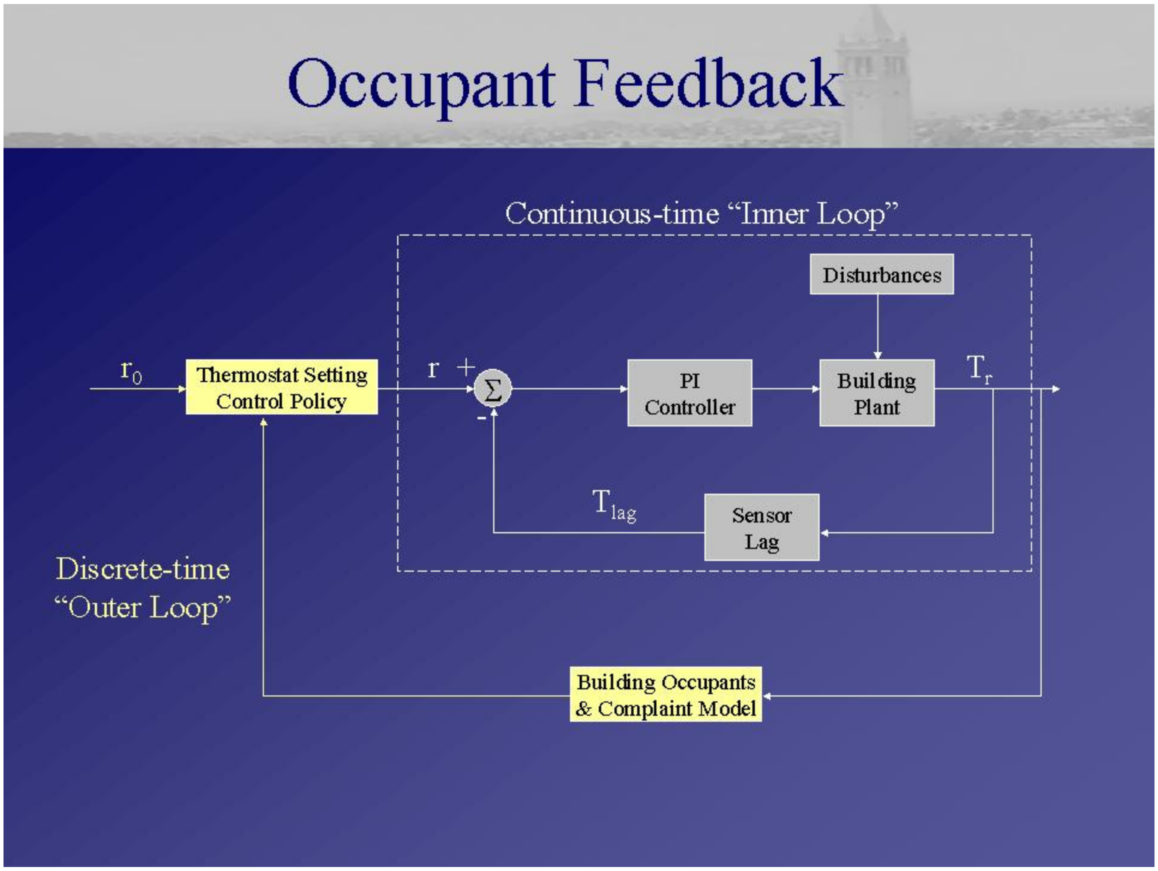

Figure 7 - Outer Loop Controller

Differentiation between the inner & outer loop control strategies is apparent in Fig. 7. The inner loop control strategy includes continuous time components, and the outer loop control strategy includes discrete time event-based setpoint control. The development of systems that use building occupants' comfort level and computational control logic in the outer loop needs extensive research and testing. Specifically, the feasibility of such a system needs to be determined by measuring the increased performance achieved by adding this new and unpredictable element to the control system.

3 Modeling & Simulation Methodology

3.1 Simulation of Complaint Behavior/Complaint Model

The simulation of complaints is based on real building occupant behavioral data characterizing the mean temperatures and frequencies at which they are generated. Obviously there are several physiological and psychological variables that determine the thermal sensation thresholds of building occupants. Reasons may range from what the person is wearing that day, what their diet is like, their general physical build, and gender, to what their mood is like that day or what their work productivity level is like. Extensive research including system identification is needed in order to develop an accurate model of complaint behavior. However, in order to simplify the simulation of the temperatures at which building occupants complain, estimated statistical parameters regarding complaint temperatures and associated rates of change from Federspiel ("Predicting", 1998, p. 7) can be used. The level-crossing model of complaint behavior as described in Sec. 2.3.2 requires that the high-temperature and low-temperature levels are the output of a system with a relative degree of at least two. This is necessary to allow for all estimated statistical parameters to be used in developing a simple model with unknown system parameters. Therefore the statistical parameters defining the output of such a system will be used to derive these unknown system parameters. The systems generating these levels are second-order "coloring filters", whose input is white Gaussian noise. Because the input noise is not zero mean, the transfer functions for both hot and cold coloring filters must have unity gain, as follows:

$$\frac{T}{n} = \frac{\omega^2}{s^2 + 2\zeta\omega s + \omega^2}, \ n \sim N(\mu_T, \sigma_n^2)$$

Equation 1 - Complaint Process Coloring Filter

Where: T = Temperature at which building occupants complain
ω = natural frequency of complaint process
ζ = damping ratio of complaint process
n = white Gaussian input noise to coloring filter
$N(\mu_T,\sigma_n^2)$ = white Gaussian noise with mean μ_T, and variance σ_n^2

The only known statistical parameter above is μ_T. Therefore more data will be required to solve for the unknown system parameters ω, and ζ. Simulating input conditions will require normalizing a zero mean white Gaussian input noise probability density function with unity variance, N(0,1), as follows:

$$N(m_x, \sigma_x^2) = \sigma_x N(0,1) + m_x$$

Equation 2 - Non-zero mean white Gaussian with non-unity variance

This is necessary due to the version of MATLAB® used, which only provides a zero-mean normal distribution with unity variance. The relationship between the values of ω, ζ and σ_n^2 can be derived from tables (Newland, 1984, App. 1), or from using the steady-state solution to the Lyapunov equation involving the covariance matrix (Tomizuka, ME233, 1999, pp. PR-6, 7) as follows:

$$\mathbf{A_c X_{ss}} + \mathbf{X_{ss} A_c^T} = -\mathbf{B_c} Q_c \mathbf{B_c^T}$$

Equation 3 - Continuous Lyapunov Equation

where $Q_c = \sigma_n^2$. In order to solve this equation, a state-space representation of the transfer function in Eqn. 1 must be obtained, as follows:

$$\dot{\mathbf{x}}(t) = \mathbf{A_c x}(t) + \mathbf{B_c} n(t) \qquad y(t) = \mathbf{C_c x}(t)$$

$$\text{where } \mathbf{A_c} = \begin{bmatrix} 0 & 1 \\ -\omega^2 & -2\zeta\omega \end{bmatrix}, \quad \mathbf{B_c} = \begin{bmatrix} 0 \\ \omega^2 \end{bmatrix}, \text{ and } \mathbf{C_c} = [1 \ 0]$$

$$\text{Define } \mathbf{x}(t) \equiv \begin{bmatrix} T(t) \\ \dot{T}(t) \end{bmatrix}, \text{ with Initial Conditions}: \mathbf{x}_0 \equiv \begin{bmatrix} T_0 \\ \dot{T}_0 \end{bmatrix} = \begin{bmatrix} \sigma_T N(0,1) + \mu_T \\ \sigma_{\dot{T}} N(0,1) \end{bmatrix}$$

Equation 4 - State-Space Realization of Hot Complaint Process Coloring Filter

The mean μ_T, and standard deviations σ_T, and $\sigma_{\dot{T}}$, are given by Federspiel's empirical data ("Predicting", 1998, p. 7), for both hot & cold levels. The solution to Eqn. 3, $\mathbf{X}_{ss}$, can be solved for analytically using the fact that it will be symmetric. This steady-state covariance matrix can be written in terms of the unknown system parameters ω, ζ, and the input noise covariance Q_c:

$$\mathbf{X}_{ss} = \begin{bmatrix} \frac{\omega}{4\zeta} Q_c & 0 \\ 0 & \frac{\omega^3}{4\zeta} Q_c \end{bmatrix}$$

Equation 5 - Solution to Continuous Lyapunov Equation

$\mathbf{X}_{ss}$ is the steady-state version of the covariance matrix, $\mathbf{X}_{xx}$. Using the fact that $y=\mathbf{C_c x}(t)$, the output covariance, X_{yy} can be computed:

$$X_{yy} = E[y(t)y^T(t)] = E[\mathbf{C_c x}(t)\mathbf{x^T}(t)\mathbf{C_c^T}] = \mathbf{C_c}E[\mathbf{x}(t)\mathbf{x^T}(t)]\mathbf{C_c^T} = \mathbf{C_c X_{xx} C_c^T}$$

E[·] is the expectation operator. The steady-state output covariance is therefore $X_{yy} = \mathbf{C_c X_{ss} C_c^T}$. This is equivalent to the variance of the complaint temperature, σ_T^2 given by Federspiel's empirical data ("Predicting", 1998, p. 7), giving $X_{yy} = \sigma_T^2 = \mathbf{C_c X_{ss} C_c^T}$. Using Eqn. 5 & $\mathbf{C_c}$, the final relationship among the values of ω, ζ and σ_T can be obtained:

$$\sigma_T^2 = \frac{\omega}{4\zeta} Q_c$$

Equation 6 - Variance of Complaint Temperature

Now there are three unknowns (ω, ζ, Q_c) , and one equation. Assume now that the output of the coloring filter is the rate of change of the complaint level instead of the temperature itself. Hence the transfer function changes from Eqn. 1 to the following:

$$\frac{\dot{T}}{n} = \frac{\omega^2 s}{s^2 + 2\zeta\omega s + \omega^2}, \ n \sim N(\mu_T, \sigma_n^2)$$

Equation 7 – Complaint Rate Process Coloring Filter

The only difference in the state-space realization with the new transfer function in Eqn. 7 will be that $\mathbf{C_c}$=[0 1]. This is because only the numerator has changed from Eqn. 1, not the characteristic equation. The steady-state solution to the Lyapunov equation resulting in the covariance matrix, $\mathbf{X_{ss}}$, will not change either because it is based only upon $\mathbf{A_c}$ & $\mathbf{B_c}$, not $\mathbf{C_c}$. The output of the transfer function in Eqn. 7 is given by $\dot{y}(t) = \mathbf{C_c x}(t)$. Using this fact and a similar procedure as before, the output covariance can be computed: $\mathbf{X}_{\dot{y}\dot{y}} = \sigma_{\dot{T}}^2 = \mathbf{C}_c \mathbf{X_{ss}} \mathbf{C}_c^T$. Therefore, using Eqn. 5 again & the new value for $\mathbf{C_c}$, the final relationship among the values of ω, ζ and $\sigma_{\dot{T}}$ can be obtained:

$$\sigma_{\dot{T}}^2 = \frac{\omega^3}{4\zeta} Q_c$$

Equation 8 - Variance of Rate of Change of Complaint Temperature

Now there are two equations and three unknowns. Combining Eqns. 6 and 8 by eliminating Q_c, the following relation can be obtained:

$$\sigma_{\dot{T}}^2 = \omega^2 \sigma_T^2$$

Equation 9 - Relationship Among Complaint Process Parameters

The problem has been algebraically reduced to one equation and one unknown. All but one of the parameters in Eqn. 9 is unknown, due to the empirical data available from Federspiel ("Predicting", 1998, p. 7). Therefore it can be used to find the natural frequency of complaint processes, ω, which is the unknown parameter. The free parameter remaining which canceled out of the derivation of Eqn. 9 can be chosen heuristically to allow for critical damping, such that ζ = 1. The following table summarizes the results:

	Hot Level	**Cold Level**
Mean (μ) – Known	91.0 ºF	54.5 ºF
Standard Deviation (σ) – Known	4.24 ºF	4.39 ºF
Rate of Change of Standard Deviation ($\sigma_{\dot{T}}$) – Known	0.84 ºF/hour	3.69 ºF/hour
Natural Frequency (ω) – Unknown	0.20 rad/hour	0.84 rad/hour
Damping Ratio(ζ) - Unknown	1	1

Table 1 - Empirically-Based Statistical Parameters Defining Complaint Levels

The value for Q_c is still unknown, however because numerical simulation is performed by discretizing the continuous dynamics using a zero-order hold, this continuous input noise is not a necessary piece of information. The discrete analog, Q_d, **is** required for simulation. Zero-order hold discretization proceeds according to the following formula with the sampling interval T:

$$\mathbf{x}_{k+1} = \underbrace{\mathbf{e}^{\mathbf{A_c T}}}_{\mathbf{A_d}} \mathbf{x}_k + \underbrace{\left(\mathbf{e}^{\mathbf{A_c T}} - \mathbf{I}\right)\mathbf{A_c^{-1} B_c}}_{\mathbf{B_d}} n_k, \quad n_k \sim N\left(\mu_T, \sigma_d^2\right)$$

$$\mathbf{y_k = C_d\, x_k,\ C_d = C_c}$$

Equation 10 - Discretization of Continuous-Time Coloring Filter

Above, $Q_d = \sigma_d^2$ is the discrete input noise matrix, which is the final system parameter required to perform discrete-time numerical simulation. In order to compute it, the steady-state solution to the discrete Lyapunov equation needs to be solved (Gelb et al., 1974, p. 76) as follows:

$$\mathbf{X_{ss}} = \mathbf{A_d X_{ss} A_d^T} + \mathbf{B_d} Q_d \mathbf{B_d^T}$$

Equation 11 - Discrete Lyapunov Equation

The solution, $\mathbf{X_{ss}}$, can be computed analytically again using the fact that it will be symmetric. However, because Q_d is the final parameter being sought after, not the solution itself, Eqn. 11 can be rewritten as follows:

$$\text{Let } \tilde{\mathbf{X}}_{\mathbf{ss}} Q_d = \mathbf{X_{ss}} \Rightarrow \tilde{\mathbf{X}}_{\mathbf{ss}} = \mathbf{A_d} \tilde{\mathbf{X}}_{\mathbf{ss}} \mathbf{A_d^T} + \mathbf{B_d B_d^T}$$

Because the problem has been cast into a form that is purely numerical with no unknown symbolic parameters, it can be performed numerically using the appropriate MATLAB® command (dlyap) that solves the discrete Lyapunov equation (Eqn. 11), rather than computing the solution analytically. The continuous time output covariance results can be used for discrete time:

$$\mathbf{X}_{yy} = \sigma_T^2 = \mathbf{C}_{d_1}\tilde{\mathbf{X}}_{\mathbf{ss}}\mathbf{C}_{d_1}^T \text{, and } \mathbf{X}_{\dot{y}\dot{y}} = \sigma_{\dot{T}}^2 = \mathbf{C}_{d_2}\tilde{\mathbf{X}}_{\mathbf{ss}}\mathbf{C}_{d_2}^T$$

$$\text{where } C_{d_1} = \begin{bmatrix} 1 & 0 \end{bmatrix} \text{, and } C_{d_2} = \begin{bmatrix} 0 & 1 \end{bmatrix}$$

Equation 12 – Discrete Output Covariance Equations

With knowledge of the fact that $\tilde{\mathbf{X}}_{\mathbf{ss}} Q_d = \mathbf{X}_{\mathbf{ss}}$, and Eqn. 12, two closed-form solutions for Q_d can be obtained, each in terms of the known empirical data:

$$Q_d = \frac{\sigma_T^2}{C_{d_1}\tilde{X}_{ss}C_{d_1}^T} = \frac{\sigma_{\dot{T}}^2}{C_{d_2}\tilde{X}_{ss}C_{d_2}^T}$$

Because $Q_d = \sigma_d^2$, the following equations also hold:

$$\sigma_d = \frac{\sigma_T}{\sqrt{C_{d_1}\tilde{X}_{ss}C_{d_1}^T}} = \frac{\sigma_{\dot{T}}}{\sqrt{C_{d_2}\tilde{X}_{ss}C_{d_2}^T}}$$

Equation 13 – Discrete Input Noise Statistics

The final parameter required to perform numerical simulation of the complaint levels can be computed according to the closed-form equations shown above. However, the computation can be based on either σ_T or $\sigma_{\dot{T}}$. When computing the final value for Q_d or σ_d using both methods, the answers diverge with increased sampling interval length. This can be illustrated by plotting percentage difference between the results obtained for σ_d using the different methods against the sampling interval:

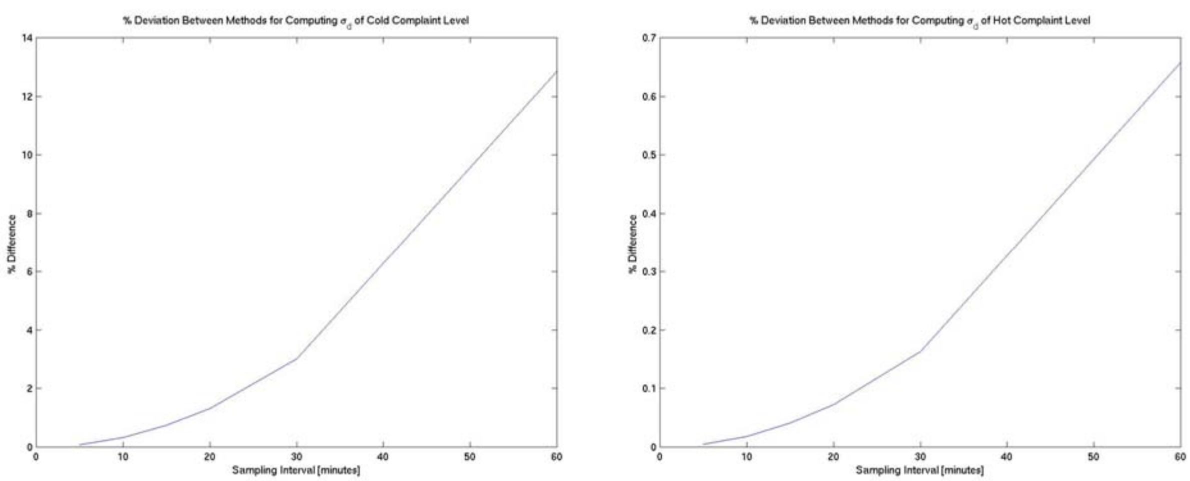

Figure 8 - % Difference Between Methods for Computing σ_d of Complaint Levels

Because σ_T, not $\sigma_{\dot{T}}$ relates to simulation of complaint levels, it will be used in the denominator of the percentage difference formula as the actual value. Therefore, even though the computation of Q_d or σ_d can be based on either method, using σ_T is the best choice. The graphs depicted in Fig. 8 illustrate that using a smaller sampling interval yields a smaller percentage difference between using either method. Hence it appears that smaller sampling intervals are advantageous due to having the option of using different analytical methods resulting in answers that are within 0.1 % error of each other. Note that the hot complaint level has a lower divergence in general than the cold complaint level.

Using Eqn. 3, the final formula to be used as the discrete input noise matrix in simulation of the complaint levels is as follows:

$$n_k = \sigma_d N(0,1) + \mu_T$$

Equation 14 - Discretized White Noise as Input to Coloring Filter

One condition needs to be met for the output covariance results of the complaint levels to exhibit the statistical characteristics expected. The simulation must run for a long enough period to allow for steady-state stationarity of the random process to be achieved. Therefore the use of the continuous & discrete steady-state Lyapunov Eqns. 3 & 11 used in computing the unknown system parameters will be valid.

3.2 Building/Room Temperature Model

3.2.1 Developing Building Plant Model

The building plant model will be kept simple by disregarding potentially complex nonlinearities introduced by modeling pressure dynamics of the fans, room and ventilation system. Doing so will provide for an unencumbered initial run at modeling these already intricate interactions between building temperature, hot & cold complaint processes, while examining possible designs for a setpoint control strategy using offline optimization. Therefore, a single zone will be used as the room model, which could be an office space divided by partitions. This space consists of 1500 square feet, with a standard 9-ft ceiling. The exact dimensions are (width = w = 30 ft) x (length = l = 50 ft) x (height = h = 9 ft). The room therefore occupies a volume of V_{air} = lwh = 13500 ft^3, and the total surface area inside of the room is A_{win} = 2wh+2wl +2lh =4440 ft^2 . It is assumed that the one wall of the room facing outside is the shorter face, such that A_{wout} = wh = 270 ft^2. The thickness of the wall on all sides is: t = 1 ft. All dimensions will be converted to metric in simulation.

A cross-sectional view of the room from the top, showing the net heat transfer modeled is as follows:

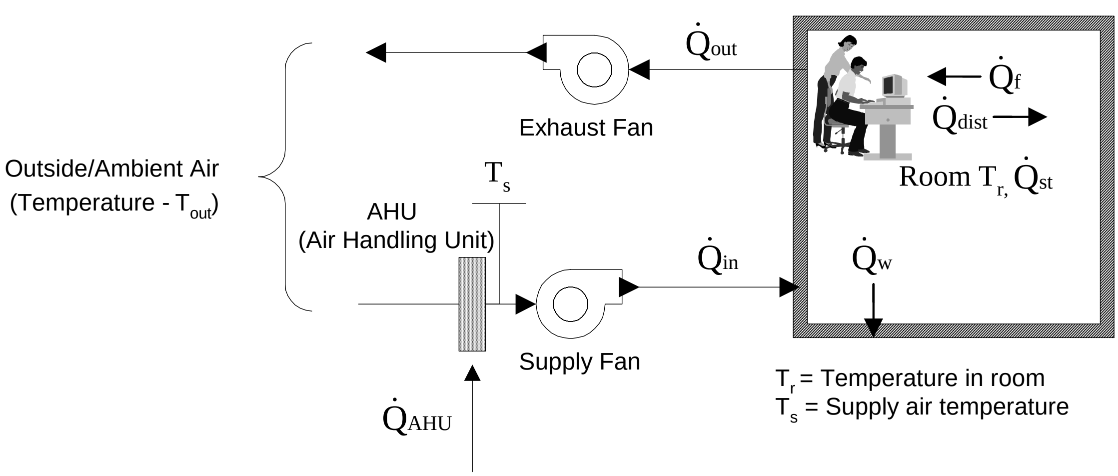

Figure 9 - Bird's Eye View of Room Heat Transfer

The following is an energy-balance equation used in modeling the heat transfer shown in Figure 9:

$$\dot{Q}_{st} = \dot{Q}_{in} - \dot{Q}_{out} + \dot{Q}_{dist} - \dot{Q}_{f} - \dot{Q}_{w}$$

Equation 15 - Energy Balance

The first state in the plant dynamics is given by the room air thermal mass storage term, $\dot{Q}_{st}$. This is given by the following equation:

$$\dot{Q}_{st} = M_r C_p \frac{dT_r}{dt}$$

Equation 16 - Room air thermal mass dynamics (1st State)

The remaining terms in the energy balance equation have the following definitions:

$$\dot{Q}_{in} = \dot{m} C_p T_s$$

A post-dimensional analysis check requires that for all temperatures to be recorded in °F, the revised equation be expressed as:

$$\dot{Q}_{in} = \tfrac{9}{5} \dot{m} C_p T_s$$

Equation 17 - Heat input from supply ventilation duct

Similarly, $\dot{Q}_{out} = \tfrac{9}{5} \dot{m} C_p T_r$

Equation 18 - Heat loss through exhaust ventilation duct

An energy balance on the heat exchange associated with the air-handling unit located upstream of the supply fan providing heating & cooling of the air in the room is as follows:

$$\dot{Q}_{AHU} = \dot{m} C_p (T_s - T_{out})$$

Again, a post-dimensional analysis check reveals that for all temperatures to be recorded in °F, the revised equation must be expressed as:

$$\dot{Q}_{AHU} = \tfrac{9}{5} \dot{m} C_p (T_s - T_{out})$$

Equation 19 - Air Handling Unit Heat Exchange Equation

The supply temperature, T_s, is not a state variable, control input, or a disturbance. It is purely an internal variable used for the convenience of describing the plant dynamics. Therefore, it may be eliminated by combining Eqns. 17 and 19. After doing so, the resulting equation representing the heat input from the supply ventilation is:

$$\dot{Q}_{in} = \tfrac{9}{5} \dot{m} C_p T_{out} + \dot{Q}_{AHU}$$

Equation 20 - Revised Heat Input from Supply Ventilation Duct

$\dot{Q}_{dist}$ is the internal heat generation, a disturbance due to building occupants and equipment, such as computers. This will be discussed in more detail later.

The first state equation/energy balance written with the parameters discussed so far, can be expressed as a combination of equations 15, 16, 18, & 20:

$$M_r C_p \frac{dT_r}{dt} = \frac{9}{5} \dot{m} C_p (T_{out} - T_r) + \dot{Q}_{AHU} + \dot{Q}_{dist} - \dot{Q}_f - \dot{Q}_w$$

Equation 21 - Room Temperature Energy Balance

The 2nd state equation can be derived from looking at the dynamics of convective heat transfer associated with the furniture in the room :

$$\dot{Q}_f = h_f A_f (T_r - T_f), \ \dot{Q}_f = M_f C_f \frac{dT_f}{dt}$$

$$\rightarrow \ M_f C_f \frac{dT_f}{dt} = h_f A_f (T_r - T_f)$$

Equation 22 - Furniture Temperature Energy Balance

The remaining states are associated with the wall. It has been found that for a wall, the simplest modeling representation is 2nd – order using lumped resistive & capacative elements, as follows (Seem, 1987, pp. 20-22) :

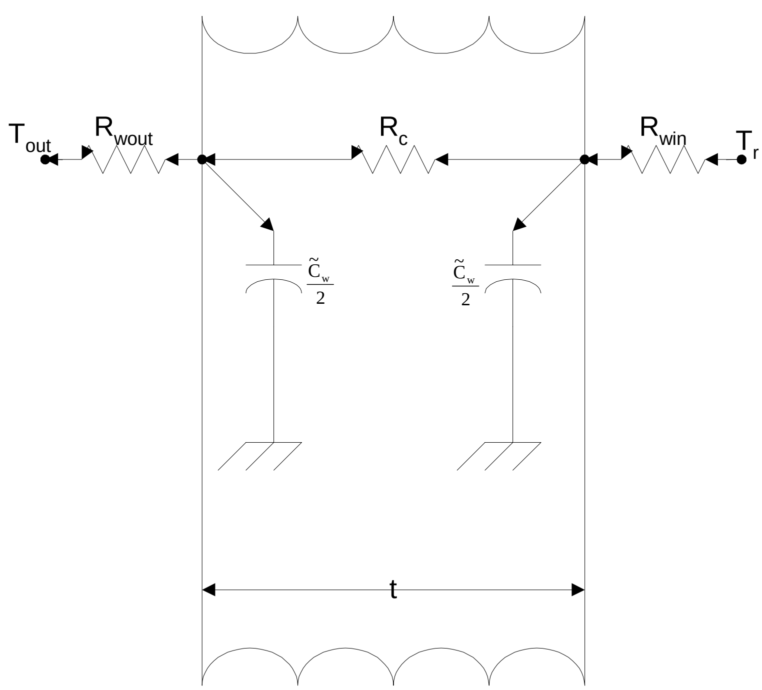

Figure 10 - Seem's representation of a wall into lumped elements

The 3rd & 4th state equations associated with the wall can be derived, using the following lumped parameters :

$$R_{w_{out}} = \frac{1}{h_{w_{out}} A_{w_{out}}} \qquad R_{w_{in}} = \frac{1}{h_{w_{in}} A_{w_{out}}} \qquad R_c = \frac{t}{kA_{w_{in}}} \quad , \quad \tilde{C}_w = M_w C_w$$

Equation 23 - Lumped Wall Parameters

Now fundamental heat flow & energy balance principles can be applied to the illustration in Fig. 10:

$\dot{Q}_{c_{w_1}} = \dot{Q}_{R_{w_{in}}} - \dot{Q}_{R_c}$ and $\dot{Q}_{c_{w_2}} = \dot{Q}_{R_c} - \dot{Q}_{R_{w_{out}}}$, where the two states associated with the wall are :

$\dot{Q}_{c_{w_1}} = \frac{\tilde{C}_w}{2}\frac{dT_{w_{in}}}{dt} = \frac{M_w C_w}{2}\frac{dT_{w_{in}}}{dt}$, and $\dot{Q}_{c_{w_2}} = \frac{\tilde{C}_w}{2}\frac{dT_{w_{out}}}{dt} = \frac{M_w C_w}{2}\frac{dT_{w_{out}}}{dt}$.

These are the heat flows through the capacative lumped elements shown in Fig. 10.

Note: $\dot{Q}_w = \dot{Q}_{R_{w_{in}}}$ → This gives $\dot{Q}_w$ from the state Eqn. 21. Additionally,

$$\dot{Q}_{R_{w_{in}}} = \frac{1}{R_{w_{in}}}(T_r - T_{w_{in}}) = h_{w_{in}} A_{w_{in}} (T_r - T_{w_{in}})$$
$$\dot{Q}_{R_c} = \frac{1}{R_c}(T_{w_{in}} - T_{w_{out}}) = \frac{kA_{w_{in}}}{t}(T_{w_{in}} - T_{w_{out}})$$
$$\dot{Q}_{R_{w_{out}}} = \frac{1}{R_{w_{out}}}(T_{w_{out}} - T_{out}) = h_{w_{out}} A_{w_{out}} (T_{w_{out}} - T_{out})$$

Combining all of the above equations, the two state equations associated with the wall are:

$$\frac{M_w C_w}{2}\frac{dT_{w_{in}}}{dt} = h_{w_{in}} A_{w_{in}} (T_r - T_{w_{in}}) - \frac{kA_{w_{in}}}{t}(T_{w_{in}} - T_{w_{out}})$$

and $$\frac{M_w C_w}{2}\frac{dT_{w_{out}}}{dt} = \frac{kA_{w_{in}}}{t}(T_{w_{in}} - T_{w_{out}}) - h_{w_{out}} A_{w_{out}} (T_{w_{out}} - T_{out})$$

Equation 24 - Derivation of Wall State Equations

A list of all state equations using the basic wall model represented in Fig. 10, and combining Eqns. 21, 22, and 24 is as follows:

1) $$M_r C_p \frac{dT_r}{dt} = \frac{9}{5}\dot{m}C_p (T_{out} - T_r) - h_f A_f (T_r - T_f) - h_{w_{in}} A_{w_{in}} (T_r - T_{w_{in}}) + \dot{Q}_{AHU} + \dot{Q}_{dist}$$

2) $$M_f C_f \frac{dT_f}{dt} = h_f A_f (T_r - T_f)$$

3) $$\frac{M_w C_w}{2}\frac{dT_{w_{in}}}{dt} = h_{w_{in}} A_{w_{in}} (T_r - T_{w_{in}}) - \frac{kA_{w_{in}}}{t}(T_{w_{in}} - T_{w_{out}})$$

4) $$\frac{M_w C_w}{2}\frac{dT_{w_{out}}}{dt} = \frac{kA_{w\,in}}{t}(T_{w_{in}} - T_{w_{out}}) - h_{w_{out}} A_{w_{out}}(T_{w_{out}} - T_{out})$$

Equation 25 - All State Equations Based on 2nd Order Wall

Now that the state equations have been established, the constants used in them can be defined:

R_{wout} = effective resistance of exterior wall surface convective heat transfer coefficient
R_{win} = effective resistance of interior wall surface convective heat transfer coefficient
R_c = effective resistance of thermal conductivity of wall
k = thermal conductivity of wall
t = thickness of wall

as well as some new ones:

l = length of room
w = width of room
h = height of room
t = thickness of wall

M_r = mass of air in room = $\rho_{air}V_{air}$
M_w = mass of wall = $\rho_{wall} V_{wall}$
M_f = mass of furniture = $m_{furn}N_p$

ρ_{air} = density of room air, V_{air} = volume of room air = lwh
ρ_{wall} = density of wall, V_{wall} = volume of wall = $A_{win}t$
A_{win}= total surface area = 2wh+2wl +2lh
m_{furn} = mass of furniture allowed per person
N_p = effective number of people in the space = A_{fl}/A_p
A_{fl} = square area of floor space = wl, A_p = square area allocated per person
A_{wout} = surface area of room facing outside = wh
$A_f = N_p(DFA_{face}+PA_{part})$, D = number of desks allowed per person, F = number of faces per desk,
P = number of partitions allowed per person
A_{face} = average square area per face of desk, A_{part} = average square area of a partition

C_p = heat capacity of room air, C_w = heat capacity of wall, C_f = heat capacity of furniture

h_{win} = inside/wall surface heat transfer coefficient
h_{wout} = exterior/wall surface heat transfer coefficient, h_f = furniture heat transfer coefficient

$\dot{m}$ = air flow rate in room = $\rho_{air}Q_{air}$, where Q_{air} = volumetric flow rate of air in room

Numerical Values of Constants Used:

$\left(\frac{Q_{air}}{A_{fl}}\right) \cong 0.15 \frac{\text{cfm}}{\text{sf of floor space}}$ is the standard room airflow rate used in buildings.

In order to obtain Q_{air}, multiply by A_{fl} and convert to metric, giving Q_{air} = 0.1062 m^3/sec.

h_{win} = inside/wall surface convective heat transfer coefficient = $h_c+5.72\varepsilon \frac{W}{m^2K}$
where ε = surface emissivity, and h_c is dependent on the direction of heat flow.

Therefore, since the ceiling and the floor are horizontal, and the four walls are vertical, the corresponding values for h_c are to be used: (Rohsenow, 1985, pp. 9-13,14)

horizontal h_c=3.08 (ceiling & floor)
upward h_c=4.04
downward h_c=0.92
The upward & downward coefficients correspond to the four walls.

Therefore, $h_c=2.68\frac{W}{m^2K}$, by taking the average over all six faces of the room.

Let the surface emissivity, ε=0.9 for the room walls, so that $h_{win} = h_c+5.72\varepsilon = 7.83 \frac{W}{m^2K}$

The exterior wall convective surface heat transfer coefficient can be calculated by using the following formula:

$h_{wout}=aV^2+bV+c$ W/m^{2}K , where V = wind speed in mph near the outside wall (Rohsenow, 9-13,14)
Assume that $V_{avg} \approx 5$ mph
For concrete: a = 0, b =1.874, c = 10.788 (Rohsenow, 1985, pp. 9-13,14)
Therefore, h_{wout}=20.16 W/m^{2}K

C_p = 1007 J/kgK at standard atmospheric pressure & temperature (Incropera, 1990, App. A)
ρ_{air} = 0.075 lb/ft^3 = 1.2 kg/m^3 (Incropera, 1990, App. A)

Assuming that the walls are made of concrete stone mix, then
C_w = 880 J/kgK, and ρ_w = 2300 kg/m^3 (Incropera, 1990, App. A)

And the furniture is mostly made up of stainless steel, then
C_f = 900 J/kgK. And let $h_f = h_{win}$ = 7.83 W/m^{2}K (Incropera, 1990, App. A)

Then finally, to calculate M_f, the following parameters are needed:

A_{fl} = square area of floor space = wl = 1500 ft^2, A_p = square area allocated per person = 300 ft^2
D = number of desks allowed per person = 1
F = number of faces per desk = 2.5
P = number of partitions allowed per person = 1
A_{face} = average square area per face of desk = 2 m^2
A_{part} = average square area of a partition = 8 m^2

Now that all constants and parameters have been defined, a post-dimensional analysis check must be performed. Additionally, substitution of the more important parameters listed above into the state equations from Eqn. 25 is required to give the final set of state equations for this model:

1) $$\frac{dT_r}{dt} = \frac{Q_{air}}{V_r}(T_{out} - T_r) - \frac{h_f A_f}{\rho C_p V_r}(T_r - T_f) - \frac{h_{w_{in}} A_{w_{in}}}{\rho C_p V_r}(T_r - T_{w_{in}}) + \frac{5}{9\rho C_p V_r}(\dot{Q}_{AHU} + \dot{Q}_{dist})$$

2) $$\frac{dT_f}{dt} = \frac{h_f A_f}{M_f C_f}(T_r - T_f)$$

3) $$\frac{dT_{w_{in}}}{dt} = \frac{2h_{w_{in}} A_{w_{in}}}{M_w C_w}(T_r - T_{w_{in}}) - \frac{2kA_{w_{in}}}{tM_w C_w}(T_{w_{in}} - T_{w_{out}})$$

$$4)\quad \frac{dT_{w_{out}}}{dt} = \frac{2kA_{w_{in}}}{tM_wC_w}(T_{w_{in}} - T_{w_{out}}) - \frac{2h_{w_{out}}A_{w_{out}}}{M_wC_w}(T_{w_{out}} - T_{out})$$

Equation 26 - Final State Equations for 2nd order Wall Model

The first disturbance is the outside temperature, T_{out}. This disturbance is based on a file containing actual TMY weather data from Sacramento, CA. There are 8760 entries in this data file corresponding to each hour of the year. The first data point corresponds to the first hour of the year, starting January 1, at 1 am. Data is not actually taken from any particular year, however monthly data is taken at random from different years and put together to form a profile of the weather for the Sacramento area. The user will specify simulation parameters including the sampling interval in minutes, the run-time length in days, and which season of the year to begin sampling weather data from. Based upon these inputs, the weather data is loaded and manipulated to allow for starting the simulation of weather data at the beginning of the particular solstice designated by the user. If needed, the data file will "wrap-around" to the beginning of the data file if the end of the year is reached before the run-time length has finished. Furthermore, since the data file gives hourly data only, if the sampling interval is less than one hour, linear interpolation is used to calculate intermediate temperature points. The following is an example of what the weather data looks like for a run-length of 60 days using a 30 minute sampling interval, beginning in the winter:

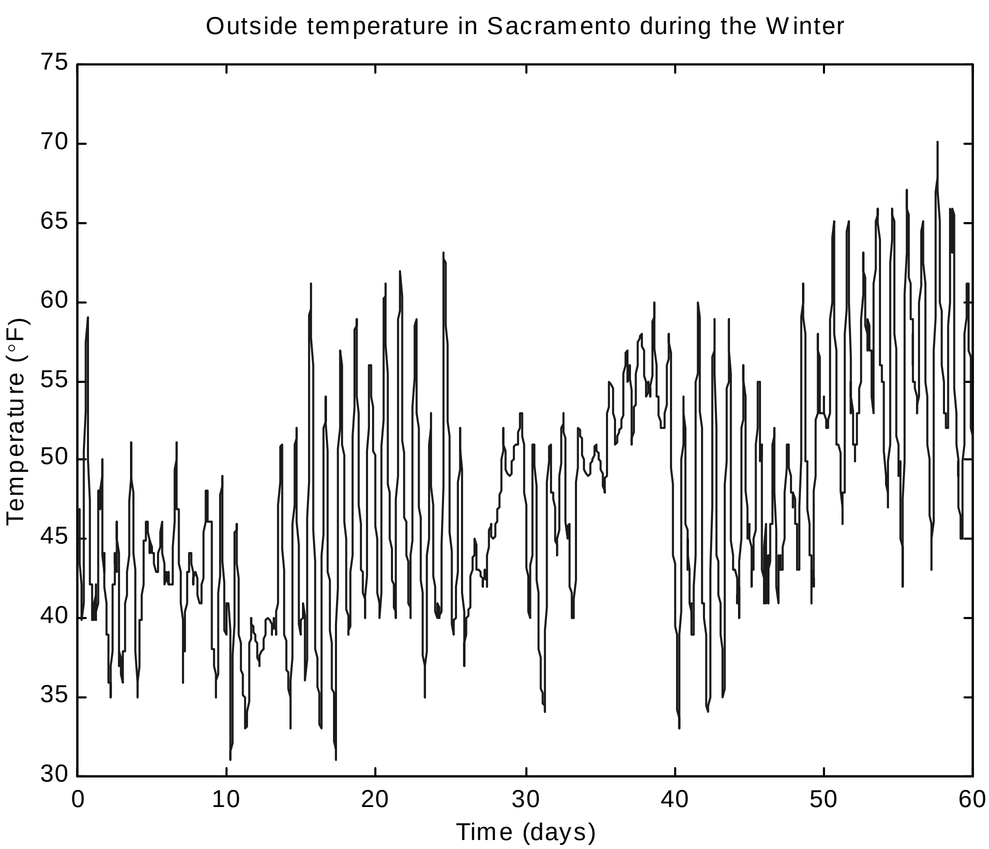

Figure 11 - Outside Temperature Disturbance

The second disturbance is internal heat generation, $\dot{Q}_{dist}$, from building occupants and equipment (computers, for example). This process is modeled as a pseudo-random binary process, in which each computer and person in the room output 140W and 100W of heat, respectively. Recall the effective number of people in the space can be calculated as follows:

N_p = effective number of people in the space = A_{fl}/A_p
A_{fl} = square area of floor space = wl, A_p = square area allocated per person

Assuming one computer is allowed per person, and it is on only when the building occupant is there, the disturbance profile for computers will be the same as for the building occupants, scaled by 1.4. The pseudo-random binary process is implemented by varying the duty cycle of a square wave randomly as the simulation progresses. This represents each person entering and leaving the space at differing times throughout the day. The % duty cycle is computed by choosing a uniformly distributed random number on the interval [0,1] and multiplying it by 100. Hence, a square wave is the best way to simulate the heat that building occupants & their computers generate. The period is based upon an 8-hour maximum length of stay. The total heat generated by all computers and building occupants in the space is summed up for a particular instant giving the total magnitude of the disturbance for that time interval. This is performed throughout the duration of the simulation. The following plot is an example of the pseudo-random binary process, over the course of a 3-day period, using a 30-min sampling interval:

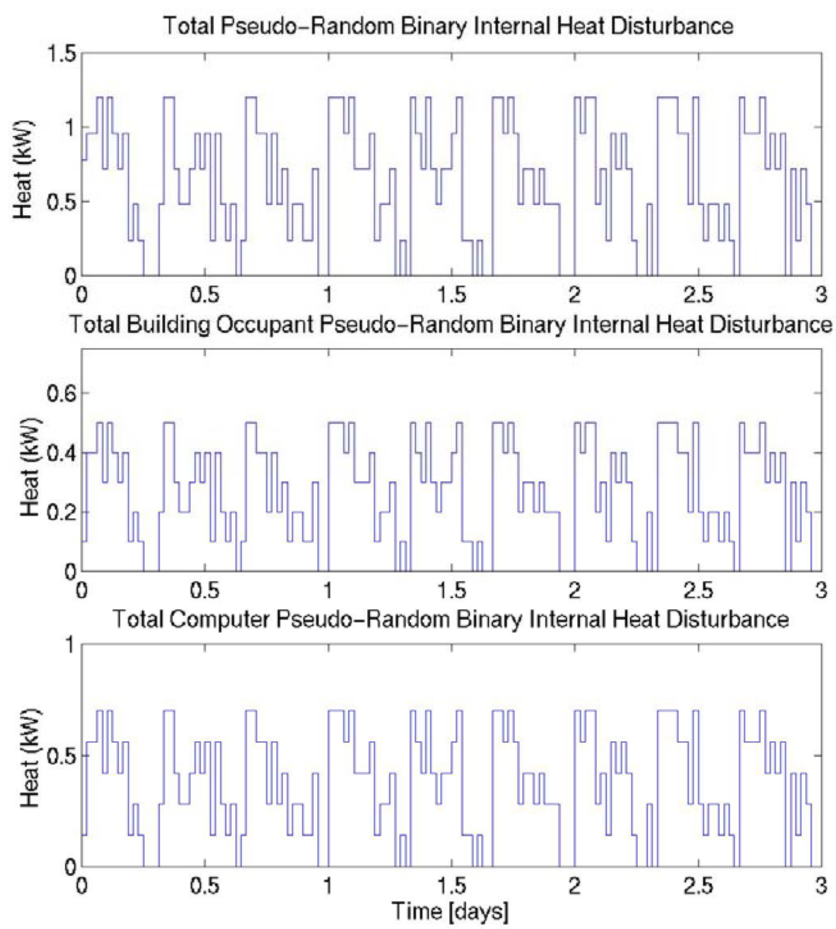

Figure 12 - Internal Heat Generation Disturbance

Although the lumped parameter method from Seem (1987, pp. 20-22) is convenient for creating a compact model, heat transfer through the wall is not necessarily best represented by two discrete states.

Continuous one-dimensional heat transfer through a wall is best described by the following governing partial differential equation for heat conduction & diffusivity through a wall slab (Incropera, 1990, p. 56):

$$\rho C_w \frac{\partial T_w}{\partial t} = k \frac{\partial^2 T_w}{\partial x^2}$$

where k = Thermal conductivity of wall
ρ = Density of wall
C_w = Heat capacity of wall
T_w = Wall temperature
x = Distance into wall
t = Time

Equation 27 - Governing Equation for Heat Conduction & Diffusivity

Using basic principles from linear algebra (Strang, 1988, pp. 52-53), a finite-difference approximation to discretize the continuous partial differential equation into matrix form across the distance into the wall can be used. The method used comes from linear graph theory (Tomizuka, ME232, 1997, pp. 40-61), hence allowing Seem's two-state wall model to be extended to any finite number of states, providing greater flexibility in approximating the governing PDE. The more finite elements contained in the wall, the higher order the model will be. The tradeoff is that the finite difference model will more closely approximate the governing PDE. In simulation, the number of finite elements to use acts as a simulation parameter, so the model remains flexible. A diagram to encapsulate the entire plant, wall, furniture, and room dynamics is a good starting point to derive & generalize the plant dynamics from a linear algebra/linear graph theory perspective:

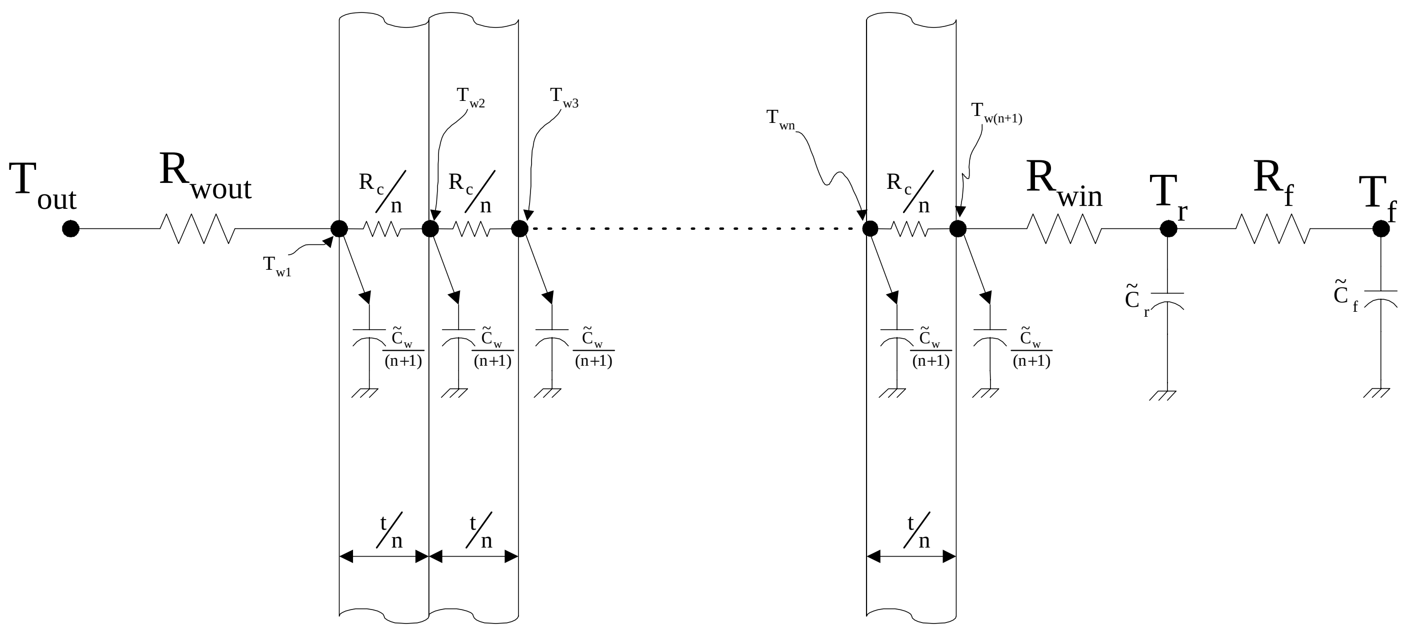

Figure 13 - Generalized Flexible Node Model

In Fig.13, n = # of finite wall elements

$$\tilde{C}_r = M_r C_p$$
$$\tilde{C}_f = M_f C_f$$
$$R_f = \frac{1}{h_f A_f}$$

Fig. 13 only represents resistive & capacative energy storage elements. To account for heat loss & gain due to the air-handling unit, building occupants, computers & ventilation, source elements need to be introduced. From linear graph theory, the system graph is as follows:

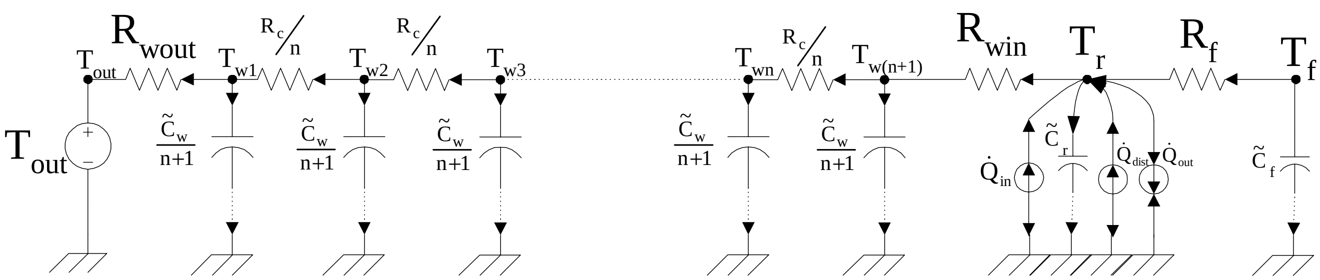

Figure 14 - Model Represented as Linear Graph

Recall: $\dot{Q}_{in} = \frac{9}{5}\dot{m}C_p T_{out} + \dot{Q}_{AHU}$

$$\dot{Q}_{out} = \frac{9}{5}\dot{m}C_p T_r$$

$\dot{Q}_{dist}$: defined as before

A standard analysis from linear graph theory is required to obtain a state-space model of the plant dynamics. The first step is to count branches, nodes, & sources:

B = # of branches = 2n+10
N = # of nodes = n+5
S = # of total sources = 4
S_A= # of across sources = 1
S_T = # of through sources = 3

The normal tree is drawn with no loops, and for this model only across sources and A-type energy storage elements exist:

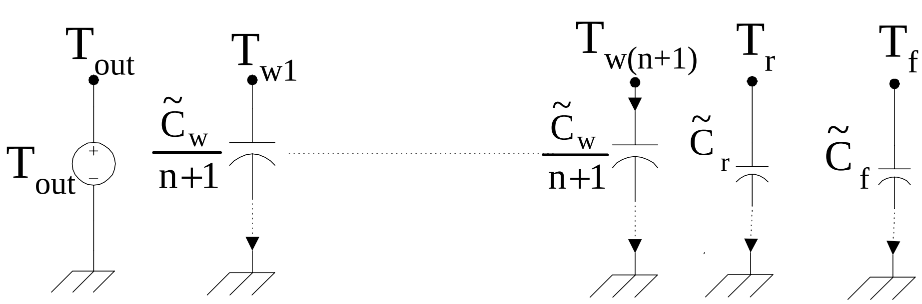

Figure 15 - Normal Tree

The states, primary & secondary variables can be derived from the normal tree:

States: $\{T_{w_1},........,T_{w_{n+1}},T_r,T_f\}$

Primary Variables: $\{T_{w_1},.........,T_{w_{n+1}},T_f,\dot{Q}_{R_{w_{out}}},\dot{Q}_{R_{c_1}},.......,\dot{Q}_{R_{c_n}},\dot{Q}_{R_{w_{in}}},\dot{Q}_{R_f}\}$

Secondary Variables: $\{\dot{Q}_{w_1},.......,\dot{Q}_{w_{n+1}},\dot{Q}_r,\dot{Q}_f,T_{R_{w_{out}}},T_{R_{c_1}},........,T_{R_{c_n}},T_{R_{w_{in}}},T_{R_f}\}$

There are B-S = 2n+6 Elemental Equations/Constitutive Relations:

$$\frac{dT_{w_1}}{dt}=\frac{n+1}{\tilde{C}_w}\dot{Q}_{w_1} \qquad \dot{Q}_{R_{w_{out}}}=\frac{1}{R_{w_{out}}}T_{R_{w_{out}}}$$

$$\vdots \qquad \dot{Q}_{R_{c_1}}=\frac{n}{R_c}T_{R_{c_1}}$$

$$\vdots \qquad \vdots$$

$$\frac{dT_{w_{n+1}}}{dt}=\frac{n+1}{\tilde{C}_w}\dot{Q}_{w_{n+1}} \qquad \vdots$$

$$\frac{dT_r}{dt}=\frac{1}{\tilde{C}_r}\dot{Q}_r \qquad \dot{Q}_{R_{c_n}}=\frac{n}{R_c}T_{R_{c_n}}$$

$$\frac{dT_f}{dt}=\frac{1}{\tilde{C}_f}\dot{Q}_f \qquad \dot{Q}_{R_{w_{in}}}=\frac{1}{R_{w_{in}}}T_{R_{w_{in}}}$$

$$\dot{Q}_{R_f}=\frac{1}{R_f}T_{R_f}$$

Performing a post-dimensional analysis check on these constitutive relationships yields the following corrections:

$$\frac{dT_{w_1}}{dt}=\frac{5(n+1)}{9\tilde{C}_w}\dot{Q}_{w_1} \qquad \dot{Q}_{R_{w_{out}}}=\frac{9}{5R_{w_{out}}}T_{R_{w_{out}}}$$

$$\vdots \qquad \dot{Q}_{R_{c_1}}=\frac{9n}{5R_c}T_{R_{c_1}}$$

$$\vdots \qquad \vdots$$

$$\frac{dT_{w_{n+1}}}{dt}=\frac{5(n+1)}{9\tilde{C}_w}\dot{Q}_{w_{n+1}} \qquad \vdots$$

$$\frac{dT_r}{dt}=\frac{5}{9\tilde{C}_r}\dot{Q}_r \qquad \dot{Q}_{R_{c_n}}=\frac{9n}{5R_c}T_{R_{c_n}}$$

$$\frac{dT_f}{dt}=\frac{5}{9\tilde{C}_f}\dot{Q}_f \qquad \dot{Q}_{R_{w_{in}}}=\frac{9}{5R_{w_{in}}}T_{R_{w_{in}}}$$

$$\dot{Q}_{R_f}=\frac{9}{5R_f}T_{R_f}$$

Equation 28 - Elemental Equations / Constituitive Relationships

There are N-1-S_A = n+3 Continuity Equations, which are the equivalent of Node Laws, or applying Kirchoff's Current Law (KCL):

$$\dot{Q}_{R_{c_1}} - \dot{Q}_{R_{w_{out}}} - \dot{Q}_{w_1} = 0$$
$$\dot{Q}_{R_{c_2}} - \dot{Q}_{R_{c_1}} - \dot{Q}_{w_2} = 0$$
$$\vdots$$
$$\vdots$$
$$\dot{Q}_{R_{c_n}} - \dot{Q}_{R_{c_{n\text{-}1}}} - \dot{Q}_{w_n} = 0$$
$$\dot{Q}_{w_{in}} - \dot{Q}_{R_{c_n}} - \dot{Q}_{w_{n+1}} = 0$$
$$\dot{Q}_{in} + \dot{Q}_{dist} - \dot{Q}_r - \dot{Q}_{R_{w_{in}}} - \dot{Q}_{R_f} - \dot{Q}_{out} = 0$$

The last equation shown above can actually be combined with Eqns. 18 & 20, to obtain the node law in terms of the variables that will actually be used to form the final state equations, giving:

$$\tfrac{9}{5}\dot{m}C_p(T_{out} - T_r) + \dot{Q}_{AHU} + \dot{Q}_{dist} - \dot{Q}_r - \dot{Q}_{R_{w_{in}}} - \dot{Q}_{R_f} = 0$$

The final continuity equation is:

$$\dot{Q}_{R_f} - \dot{Q}_f = 0$$

Equation 29 - Contunuity Equations (Node Laws using Kirchoff's Current Law - KCL)

There are $B\text{-}N\text{+}1\text{-}S_T$ = n+3 <u>Compatibility Equations</u>, which are the equivalent of <u>Loop Laws</u>, or applying <u>Kirchoff's Voltage Law</u> (KVL):

$$T_{out} - T_{w_1} + T_{R_{w_{out}}} = 0$$
$$T_{w_1} - T_{w_2} + T_{R_{c_1}} = 0$$
$$\vdots$$
$$\vdots$$
$$T_{w_n} - T_{w_{n+1}} + T_{R_{c_n}} = 0$$
$$T_{w_{n+1}} - T_r + T_{R_{w_{in}}} = 0$$
$$T_r - T_f - T_{R_f} = 0$$

Equation 30 - Compatibility Equations (Loop Laws using Kirchoff's Voltage Law - KVL)

The elemental equations in Eqn. 28 can be cast into matrix form, as follows:

$$
\frac{d}{dt}\underbrace{\begin{bmatrix} T_{w_1} \\ T_{w_2} \\ \vdots \\ \vdots \\ T_{w_{n+1}} \\ T_r \\ T_f \end{bmatrix}}_{x\,:\,n+3\text{ states}} = \underbrace{\begin{bmatrix} \frac{5(n+1)}{9\tilde{C}_w} & 0 & \cdots & \cdots & 0 & 0 & 0 \\ 0 & \ddots & & & \vdots & \vdots & \vdots \\ \vdots & & \ddots & & \vdots & \vdots & \vdots \\ \vdots & & & \ddots & 0 & 0 & 0 \\ 0 & \cdots & \cdots & 0 & \frac{5(n+1)}{9\tilde{C}_w} & 0 & 0 \\ 0 & \cdots & \cdots & 0 & 0 & \frac{5}{9\tilde{C}_r} & 0 \\ 0 & \cdots & \cdots & 0 & 0 & 0 & \frac{5}{9\tilde{C}_r} \end{bmatrix}}_{P\,:\,\text{"Capacitance" Matrix}} \underbrace{\begin{bmatrix} \dot{Q}_{w_1} \\ \dot{Q}_{w_2} \\ \vdots \\ \vdots \\ \dot{Q}_{w_{n+1}} \\ \dot{Q}_r \\ \dot{Q}_f \end{bmatrix}}_{z\,:\,\text{secondary through variables}}
$$

Equation 31 - Capacative Elemental Matrix Equation

$$
\frac{d}{dt}\underbrace{\begin{bmatrix} \dot{Q}_{R_{w_{out}}} \\ \dot{Q}_{R_{c_1}} \\ \vdots \\ \vdots \\ \dot{Q}_{R_{c_n}} \\ \dot{Q}_{R_{w_{in}}} \\ \dot{Q}_{R_f} \end{bmatrix}}_{\substack{v\,:\,n+3\text{ remaining} \\ \text{"non-state" primary} \\ \text{variables}}} = \underbrace{\begin{bmatrix} \frac{9}{5R_{w_{out}}} & 0 & \cdots & \cdots & 0 & 0 & 0 \\ 0 & \frac{9n}{5R_c} & & & \vdots & \vdots & \vdots \\ \vdots & & \ddots & & \vdots & \vdots & \vdots \\ \vdots & & & \ddots & 0 & 0 & 0 \\ 0 & \cdots & \cdots & 0 & \frac{9n}{5R_c} & 0 & 0 \\ 0 & \cdots & \cdots & 0 & 0 & \frac{9}{5R_{w_{in}}} & 0 \\ 0 & \cdots & \cdots & 0 & 0 & 0 & \frac{9}{5R_f} \end{bmatrix}}_{R\,:\,\text{"Resistive" Matrix}} \underbrace{\begin{bmatrix} T_{R_{w_{out}}} \\ T_{R_{c_1}} \\ \vdots \\ \vdots \\ T_{R_{c_n}} \\ T_{R_{w_{in}}} \\ T_{Rf} \end{bmatrix}}_{s\,:\,\text{secondary across variables}}
$$

Equation 32 - Resistive Elemental Matrix Equation

The continuity equations from Eqn. 29 can also be cast into matrix form, as follows:

$$
\underbrace{\begin{bmatrix} \dot{Q}_{w_1} \\ \dot{Q}_{w_2} \\ \vdots \\ \vdots \\ \dot{Q}_{w_{n+1}} \\ \dot{Q}_r \\ \dot{Q}_f \end{bmatrix}}_{\substack{z\,:\,\text{secondary} \\ \text{through} \\ \text{variables}}} = \underbrace{\begin{bmatrix} -1 & 1 & 0 & \cdots & \cdots & 0 & 0 \\ 0 & -1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots & \vdots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 & \vdots \\ 0 & \cdots & 0 & 0 & -1 & 1 & 0 \\ 0 & \cdots & \cdots & 0 & 0 & -1 & -1 \\ 0 & \cdots & \cdots & \cdots & 0 & 0 & 1 \end{bmatrix}}_{T\,:\,\text{"Transfer" Matrix}} \underbrace{\begin{bmatrix} \dot{Q}_{R_{w_{out}}} \\ \dot{Q}_{R_{c_1}} \\ \vdots \\ \vdots \\ \dot{Q}_{R_{c_n}} \\ \dot{Q}_{R_{w_{in}}} \\ \dot{Q}_{R_f} \end{bmatrix}}_{\substack{v\,:\,n+3\text{ remaining} \\ \text{"non-state" primary} \\ \text{variables}}} + \underbrace{\begin{bmatrix} 0 \\ 0 \\ \vdots \\ \vdots \\ 0 \\ 1 \\ 0 \end{bmatrix}}_{e_1} \underbrace{\dot{Q}_{AHU}}_{u\,:\,\text{control}} + \underbrace{\begin{bmatrix} 0 \\ 0 \\ \vdots \\ \vdots \\ 0 \\ 1 \\ 0 \end{bmatrix}}_{e_1} \underbrace{\dot{Q}_{dist}}_{w_1\,:\,\text{disturbance}} + \tfrac{9}{5}\dot{m}C_p \underbrace{\begin{bmatrix} 0 \\ 0 \\ \vdots \\ \vdots \\ 0 \\ 1 \\ 0 \end{bmatrix}}_{e_1} \underbrace{T_{out}}_{w_2\,:\,\text{disturbance}} - \tfrac{9}{5}\dot{m}C_p \underbrace{\begin{bmatrix} 0 \\ 0 \\ \vdots \\ \vdots \\ 0 \\ 1 \\ 0 \end{bmatrix}}_{e_1} \underbrace{T_r}_{e_1^T x}
$$

Equation 33 - Continuity Matrix Equation

The compatibility equations from Eqn. 30 can be cast similarly into matrix form, as follows:

$$\underbrace{\begin{bmatrix} T_{R_{w_{out}}} \\ T_{R_{c_1}} \\ \vdots \\ \vdots \\ T_{R_{c_n}} \\ T_{R_{w_{in}}} \\ T_{Rf} \end{bmatrix}}_{s:\text{secondary across variables}} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & \cdots & \cdots & 0 & 0 \\ -1 & 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & -1 & \ddots & \ddots & \ddots & \vdots & \vdots \\ 0 & \ddots & \ddots & \ddots & \ddots & 0 & \vdots \\ \vdots & \ddots & \ddots & \ddots & 1 & 0 & 0 \\ 0 & \cdots & 0 & 0 & -1 & 1 & 0 \\ 0 & \cdots & \cdots & \cdots & 0 & 1 & -1 \end{bmatrix}}_{\substack{-T^T:\text{Skew Symmetric to "Transfer" Matrix in} \\ \text{Continuity Matrix Eqn. 35}}} \underbrace{\begin{bmatrix} T_{w_1} \\ T_{w_2} \\ \vdots \\ \vdots \\ T_{w_{n+1}} \\ T_r \\ T_f \end{bmatrix}}_{x:n+3\text{ states}} - \underbrace{\begin{bmatrix} 1 \\ 0 \\ \vdots \\ \vdots \\ 0 \\ 0 \\ 0 \end{bmatrix}}_{e_2} \underbrace{T_{out}}_{w_2:\text{disturbance}}$$

Equation 34 - Compatibility Matrix Equation

The matrices used in the continuity & compatibility matrix equations are skew-symmetric to each other, almost eliminating the need for using KVL. Summarizing all of the matrix equations more compactly, (Eqns. 31 – 34):

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{Pz} \\ \mathbf{v} &= \mathbf{Rs} \\ \mathbf{z} &= \mathbf{Tv} + \mathbf{e_1 u} + \mathbf{e_1 w_1} + \tfrac{9}{5}\dot{m}C_p(\mathbf{e_1 w_2} - \mathbf{e_1 e_1^T x}) \\ \mathbf{s} &= -\mathbf{T^T x} - \mathbf{e_2 w_2} \end{aligned}$$

Elimination of extra "non-state" primary variables & all secondary variables (z,v,s) is the next step. The resulting set of state equations as a function of the states, inputs, disturbances, matrices, system parameters, & unit vectors is as follows :

$$\dot{\mathbf{x}} = \underbrace{-\mathbf{P}(\mathbf{TRT^T} + \tfrac{9}{5}\dot{m}C_p\,\mathbf{e_1 e_1^T})}_{\mathbf{A}}\mathbf{x} + \underbrace{\mathbf{Pe_1}}_{\mathbf{B}}u + \underbrace{\mathbf{Pe_1}}_{\mathbf{B_{w_1}}}w_1 + \underbrace{\mathbf{P}(\tfrac{9}{5}\dot{m}C_p\,\mathbf{e_1} - \mathbf{TRe_2})}_{\mathbf{B_{w_2}}}w_2$$

→ $$\dot{\mathbf{x}} = \mathbf{Ax} + \mathbf{B}u + \mathbf{B_{w_1}}w_1 + \mathbf{B_{w_2}}w_2$$

Let $\mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix}$, $\mathbf{B_w} = [\mathbf{B_{w_1}} \quad \mathbf{B_{w_2}}]$

Also, $y = T_r$ → $y = \mathbf{e_1^T x} = \mathbf{Cx}$ → $\mathbf{C} = \mathbf{e_1^T}$

In summary,
$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{Ax} + \mathbf{B}u + \mathbf{B_w w} \\ \mathbf{y} &= \mathbf{Cx} \end{aligned}$$

Where

$$\mathbf{A} = -\mathbf{P}(\mathbf{TRT^T} + \tfrac{9}{5}\dot{m}C_p\mathbf{e_1}\mathbf{e_1^T})$$
$$\mathbf{B} = \mathbf{Pe_1}$$
$$\mathbf{B_w} = \mathbf{P}\left[\mathbf{e_1} \quad \tfrac{9}{5}\dot{m}C_p\mathbf{e_1} - \mathbf{TRe_2}\right]$$
$$\mathbf{C} = \mathbf{e_1^T}$$

Equation 35 - Final Generalized State Equations for Plant

The equations resulting from Eqn. 35 for the case when n=1 match the equations derived from energy balance methods, validating Eqn. 26 and the generalization.

Following suit with the complaint processes described in Sec. 3.1, simulation of the model takes places by discretization of the continuous plant dynamics, using a zero-order hold. The sampling interval acts as a simulation parameter provided by the user, and a corresponding state space realization can be obtained:

$$\begin{bmatrix} \mathbf{x_{k+1}} \\ y_k \end{bmatrix} = \begin{bmatrix} \mathbf{A_d} & \mathbf{B_d} & \mathbf{B_{w_d}} \\ \mathbf{C_d} & \mathbf{D_d} & \mathbf{D_{w_d}} \end{bmatrix} \begin{bmatrix} \mathbf{x_k} \\ u_k \\ \mathbf{w_k} \end{bmatrix}$$

Equation 36 - Discretized Building Plant

3.2.1 Sensor Dynamics

At this point it is important to introduce another dynamic element of the system, prior to closing the loop with a PI controller. In practice, an intrinsic time delay is associated with the measurement of room temperature by a thermostat sensor. This lag in the sensor reading, often on the order of 7 - 8 minutes, is best represented as a first-order filter. Recall from Fig. 6 that the inner loop controller will actually involve a sensor lag element in the feedback loop, as follows:

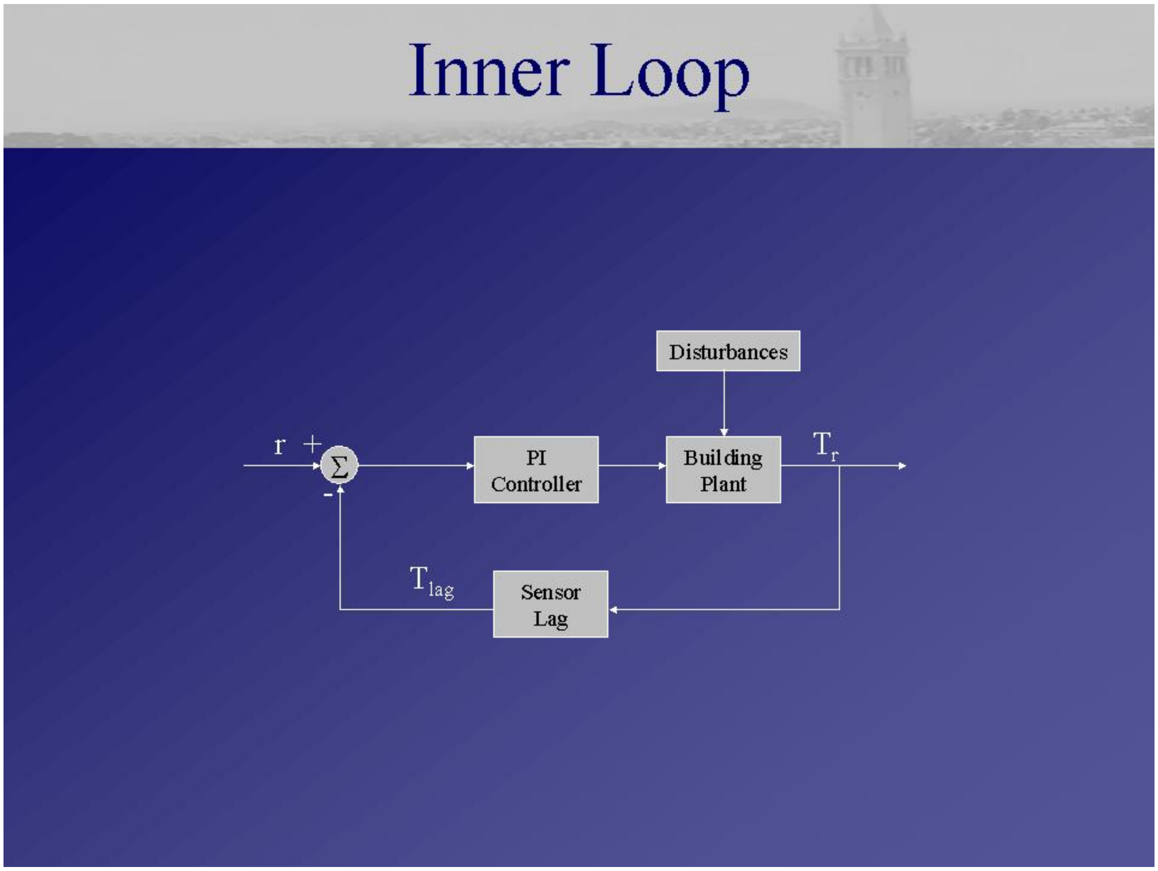

Figure 16 - Figure 6 revisited : Inner Loop Controller

The inherent sensor lag affects measurement of the room temperature prior to being used in the PI control algorithm. The transfer function of the sensor lag is as follows:

$$\frac{Y_{lag}(s)}{Y(s)} = \frac{1}{Ts + 1}$$

Where: T = Time lag in sensor
Y = Sensor measurement before lag
Y_{lag} = Sensor measurement after lag

Equation 37 - Transfer Function for Sensor Dynamics

The time lag in the sensor acts as a simulation parameter to be provided by the user. Similar to the building dynamics, the sensor dynamics must also be discretized with a zero order hold for the purposes of simulation, giving the following:

$$\begin{bmatrix} x_{lag_{k+1}} \\ y_{lag_k} \end{bmatrix} = \begin{bmatrix} A_{g_d} & B_{g_d} \\ C_{g_d} & D_{g_d} \end{bmatrix} \begin{bmatrix} x_{lag_k} \\ y_k \end{bmatrix}$$

Equation 38 - Discretized Sensor Lag

3.2.2 PI Controller

Controllers with two gains are almost always used in industry for the regulation of air space temperature. Therefore designing a simple PI (proportional-integral) controller is the best choice for comparison. Using this controller as opposed to a simpler P (proportional) controller is advantageous because the desired building temperature can be achieved without steady-state error due to the introduction of integral action. However, using higher order controllers such as the ones that might be used in optimal, LQG, or H_∞ control are not as easily implemented in practice, and cannot measure up to the PI controller's simplicity in design described by the following transfer function:

$$\frac{U(s)}{E(s)} = \frac{K_p s + K_I}{s}$$

$U(s)$ = Controller output, Plant Input
$E(s)$ = Controller input, or error which is the difference between the reference temperature and the measured room temperature
K_p = Proportional control gain
K_i = Integral control gain

Equation 39 - Transfer Function for PI controller

These PI gains act as simulation parameters to be selected by the user, and can be tuned using a three-step procedure, to be discussed shortly. Like the other system components, the PI controller will need to be realized in state space form and discretized with a zero order hold. However, the results of the discretization are needed for tuning the PI gains as well as simulation. To put the transfer function in state-space form, the first step is to cross-multiply, and take the inverse Laplace transform of both sides:

$$\dot{u}(t) = K_p \dot{e}(t) + K_i e(t)$$

Now, choosing the state: $x(t) = \int e(t)\,dt$, the state-space representation is:

$$\dot{x}(t) = 0x(t) + 1e(t)$$
$$u(t) = K_i x(t) + K_p e(t)$$

Therefore, A_c=0, B_c = 1, C_c = K_i, and D_c = K_p. The system can be discretized with sampling interval, T, using Eqn. 12:

$$A_d = e^{A_c T},\; B_d = (e^{A_c T} - I)A_c^{-1}B_c,\; C_d = C_c,\; D_d = D_c$$

Because the PI controller is a first order system, all system parameters shown above are scalar. There appears to be a problem due to the fact that A_c=0. However, B_d can be computed in spite of the apparent singularity by using L'Hospital's rule:

$$B_d = \lim_{A_c \to 0} \frac{e^{A_c T} - 1}{A_c} B_c = \lim_{A_c \to 0} T e^{A_c T} B_c = TB_c = T(1) = T$$

Therefore, A_{cd} =1, B_{cd} = T, C_{cd} = K_i, and D_{cd} = K_p, where the discrete-time state equations implementing the PI controller are as follows:

$$\begin{bmatrix} x_{c_{k+1}} \\ u_k \end{bmatrix} = \begin{bmatrix} A_{c_d} & B_{c_d} \\ C_{c_d} & D_{c_d} \end{bmatrix} \begin{bmatrix} x_{c_k} \\ e_k \end{bmatrix}$$

Equation 40 - Derivation of Discretized PI Controller

In order to tune the PI controller gains, a state-space approach or appropriate MATLAB® commands can be used. The state-space approach is conducive to preserving the states of each component of the closed-loop: building plant, sensor, and PI controller, such that they will be accessible from the resulting realization. To form the closed-loop realization, all components need to be algebraically concatenated by imposing constraints which eliminate internal variables such as the error e_k, control u_k, and output of the thermostat y_{lag_k}. Using extensive algebraic manipulation, the following closed-loop state-space representation can be derived:

$$\begin{bmatrix} \mathbf{x}_{k+1} \\ x_{c_{k+1}} \\ x_{lag_{k+1}} \\ y_k \end{bmatrix} = \begin{bmatrix} \mathbf{A_d} - \frac{\mathbf{B_d} D_{c_d} D_{g_d} \mathbf{C_d}}{1+D_d D_{c_d} D_{g_d}} & \frac{\mathbf{B_d} C_{c_d}}{1+D_d D_{c_d} D_{g_d}} & -\frac{\mathbf{B_d} D_{c_d} C_{g_d}}{1+D_d D_{c_d} D_{g_d}} & \mathbf{B_{w_d}} - \frac{\mathbf{B_d} D_{c_d} D_{g_d} \mathbf{D_{w_d}}}{1+D_d D_{c_d} D_{g_d}} & \frac{\mathbf{B_d} D_{c_d}}{1+D_d D_{c_d} D_{g_d}} \\ -\frac{B_{c_d} D_{g_d} \mathbf{C_d}}{1+D_d D_{c_d} D_{g_d}} & A_{c_d} - \frac{B_{c_d} D_{g_d} D_d C_{c_d}}{1+D_d D_{c_d} D_{g_d}} & -\frac{B_{c_d} C_{g_d}}{1+D_d D_{c_d} D_{g_d}} & -\frac{B_{c_d} D_{g_d} \mathbf{D_{w_d}}}{1+D_d D_{c_d} D_{g_d}} & \frac{B_{c_d}}{1+D_d D_{c_d} D_{g_d}} \\ \frac{B_{g_d} \mathbf{C_d}}{1+D_d D_{c_d} D_{g_d}} & \frac{B_{g_d} D_d C_{c_d}}{1+D_d D_{c_d} D_{g_d}} & A_{g_d} - \frac{B_{g_d} D_d D_{c_d} C_{g_d}}{1+D_d D_{c_d} D_{g_d}} & \frac{B_{g_d} \mathbf{D_{w_d}}}{1+D_d D_{c_d} D_{g_d}} & \frac{B_{g_d} D_d D_{c_d}}{1+D_d D_{c_d} D_{g_d}} \\ \frac{\mathbf{C_d}}{1+D_d D_{c_d} D_{g_d}} & \frac{D_d C_{c_d}}{1+D_d D_{c_d} D_{g_d}} & -\frac{D_d D_{c_d} C_{g_d}}{1+D_d D_{c_d} D_{g_d}} & \frac{\mathbf{D_{w_d}}}{1+D_d D_{c_d} D_{g_d}} & \frac{D_d D_{c_d}}{1+D_d D_{c_d} D_{g_d}} \end{bmatrix} \begin{bmatrix} \mathbf{x}_k \\ x_{c_k} \\ x_{lag_k} \\ \mathbf{w}_k \\ r_k \end{bmatrix}$$

Equation 41 – State-Space Realization of the Closed-Loop System

Conveniently, this closed-loop state-space representation is already in discrete time, because all systems have previously been discretized with a zero-order hold. Only the external inputs to the closed-loop system, w_k (disturbances) and r_k (reference) are considered inputs. Inputs to each component of the system are considered internal variables that have been absorbed into the closed-loop state-space representation. Although Eqn. 41 appears to be algebraically intimidating, comparing the closed-loop transfer function of this system with a transfer function derived from using appropriate MATLAB® commands are found to be identical. MATLAB® commands 'feedback' and 'series' are used to arrive at this transfer function, using each component of the system. Either method can be used to perform the gain tuning procedure. The closed-loop system is MISO (multi-input single-output) so it is conducive to applying a higher order more sophisticated robust H_∞ or other type of controller, as opposed to a simple PI controller. These methods can tune the closed-loop response for disturbance rejection as well as reference tracking. However, the existence of disturbances is a given for this model, and rejection of them is not of paramount importance. The PI gains need only to be based upon desired thermostat setpoint tracking. Therefore, only the single transfer function $TF_{r \to y}$ will be used in the PI gain tuning procedure. From Eqn. 41, it appears that the transfer function of interest corresponds to the third input of the closed-loop, r_k and the single output, y_k.

Each step in the PI gain-tuning procedure uses a MATLAB® optimization call based upon the Nelder-Mead simplex direct search method. It is a multidimensional unconstrained nonlinear minimization method, returning a vector that is the local minimizer of a cost function. It starts near a user-defined initial vector, and is evaluated several times during the simplex search. The cost function in question will vary depending on the step being performing, and the vector in question contains the PI gains. The three steps are as follows:

1) Optimization for stability – Due to discrete-time formulation, the eigenvalues must lie within the unit circle in the z-plane for stability. Hence, the spectral radius acts as the cost function. The initial starting vector contains the PI gains, and can be arbitrary within reason. The final resulting vector contains stabilizing PI gains and is used as the initial starting vector in the subsequent step.

2) Optimization for minimum phase behavior – In discrete time, minimum phase behavior is dictated by all of the transmission zeros of the LTI system being within the unit circle. Therefore, the absolute value of the zero furthest from the origin acts as the cost function. The gains resulting from this step are often not too different from the initial starting gains. This is due to the fact that the resulting system from step 1 is already both stable & minimum phase. If the optimizer encounters a pair of PI gains that do not yield stability during this minimum phase search, the cost function is penalized heavily by assigning a very large number as the cost. The stable minimum phase PI gains from this step are used as the initial starting gains in the final step.

3) Optimization using the minimum ISE criterion - The textbook definition (Dorf, 1998, pp. 254-256) of the ISE (integral of square error) criterion is as a performance index used as a quantitative measure to evaluate the system's performance. The ISE criterion is given by the L_2-norm, and acts as the cost function for this step:

$$\text{ISE} = \int_0^{T_s} e^2(t)dt$$

Equation 42 - Integral of Squared Error

If the optimizer returns PI gains that do not yield stability or minimum phase behavior during the search, the cost function is penalized heavily by assigning a very large number as the cost. T_s is the settling time, so the integral will approach a steady-state value. The classic ISE definition defines e(t) as the difference between a constant r and y(t), so that e(t) = r – y(t). However, specific time-series performance specifications are of interest. They are more well represented by a desired tracking trajectory as opposed to a constant, hence $y_{des}(t)$ is used instead of r. Therefore, e(t) is the error between a step response of the actual system output y(t), and the output of some desired response, $y_{des}(t)$, so that $e(t)=y(t)-y_{des}(t)$. $y_{des}(t)$ is based upon a 2nd order system with $\zeta=1$, and a user-supplied settling time. Minimization of the area between the curves defined by the step responses of the actual system output and the desired output is the goal, as both simulations progress from a zero initial condition to steady-state. Eqn. 42 must be represented in discrete time, because optimization of the ISE value will involve simulation of a discrete version of the system. An acceptable realization of this ISE metric is as follows :

$$\text{ISE} = \sum_{k=0}^{n} e_k^2 = \sum_{k=0}^{n} e_k^T e_k$$

Equation 43 - Discrete ISE criterion

The finite-valued sequence of $\{e\}_{k=0}^{n}$ can be represented as a vector, $\mathbf{e} = \{e\}_{k=0}^{n}$. Therefore, the criterion can be written more compactly: $\text{ISE} = \mathbf{e}^\mathbf{T}\mathbf{e} = \|\mathbf{e}\|_2^2$. Due to the fact that the ISE criterion yields the best response in terms of settling time and peak overshoot, it is used in lieu of other available performance indices, such as IAE (integral of absolute error), defined as :

$$\text{IAE} = \int_0^{T_s} |e(t)|dt$$

Equation 44 - Intergral of Absolute Error

In discrete form, the IAE criterion is:

$$\text{IAE} = \sum_{k=0}^{n} |e_k|$$

Equation 45 - Discrete IAE Criterion

Either version of the MATLAB® optimizer can be used when tuning the PI gains, on any operating system platform, and using either of the available methods described previously. Each different combination will yield slightly different results for the final response. The best results are found when using MATLAB®'s 'fmins' optimizer on MS Windows, and the state-space method for the final closed-loop dynamics. The step response plots resulting from applying the PI-gain tuning procedure using these methods is as follows:

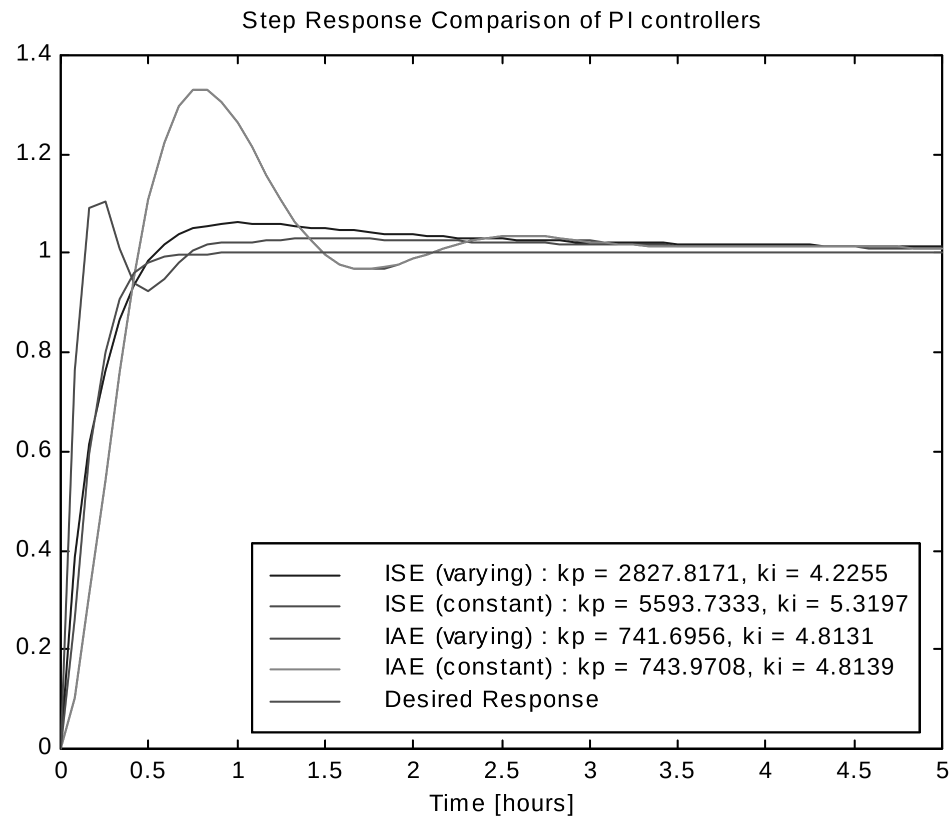

Figure 17 - Closed-Loop Step Response Using PI Controller Method # 1

For comparison, the results of applying the ISE criterion and the IAE criterion method both with constant reference and varying trajectory-tracking reference are shown. Using a different version of the MATLAB® optimizer, 'fminsearch', on a different operating system platform, UNIX, yields slightly different results for the final response after PI gain-tuning, and is as follows:

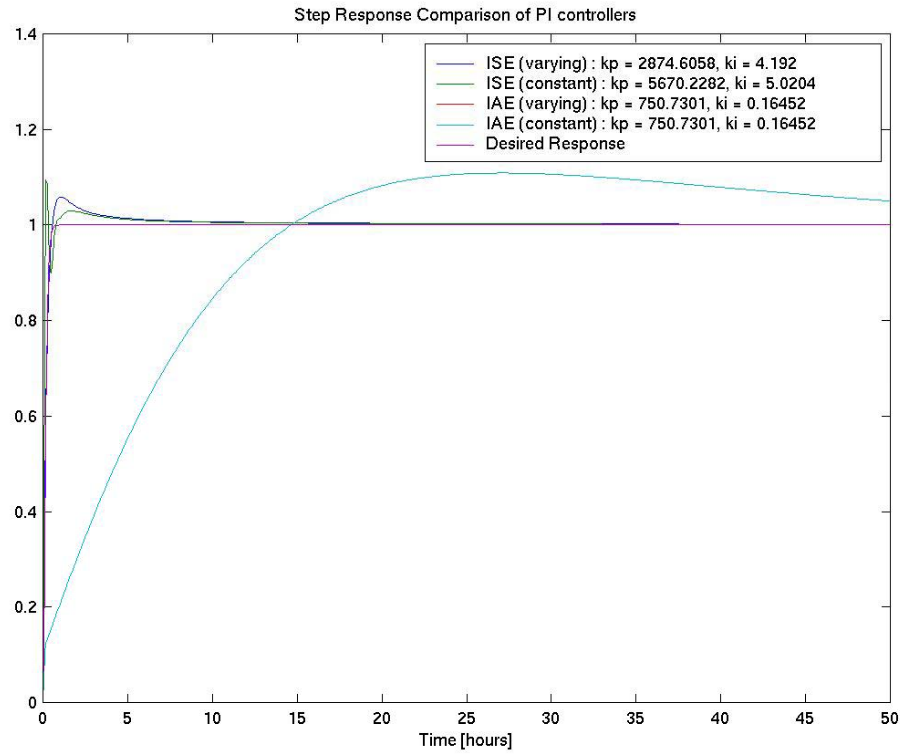

Figure 18 - Closed-Loop Step Response Using PI Controller Method # 2

In Fig. 18, the IAE criterion yields PI gains with a very sluggish response, unlike Fig 17. Therefore, there appears to be a large discrepancy in the final response when using the PI gains tuned with the IAE method, whereas the ISE method is more consistent. However, regardless of the platform, version of the optimizer, or method used to form the closed-loop, using a varying reference tracking command and the ISE method results in a superior step response. Therefore, the ISE criterion is to be used when tuning the PI gains for all simulations. The optimal PI gains tuned using this procedure also appear to be dependent on the number of finite wall elements chosen. The results of the PI gain tuning procedure using varying finite wall elements are shown in the following table & graphs:

Wall sections	K_p	K_I
0	2756.6763	4.5566
1	2811.8951	4.4906
2	2858.532	4.4155
5	2874.6058	4.192
10	3018.1166	3.9001
20	3019.7722	3.9672
50	2966.6671	3.9407
100	2943.1706	3.9221
Average	2906.1795	4.1731

Table 2 - PI Gain Tuning Results

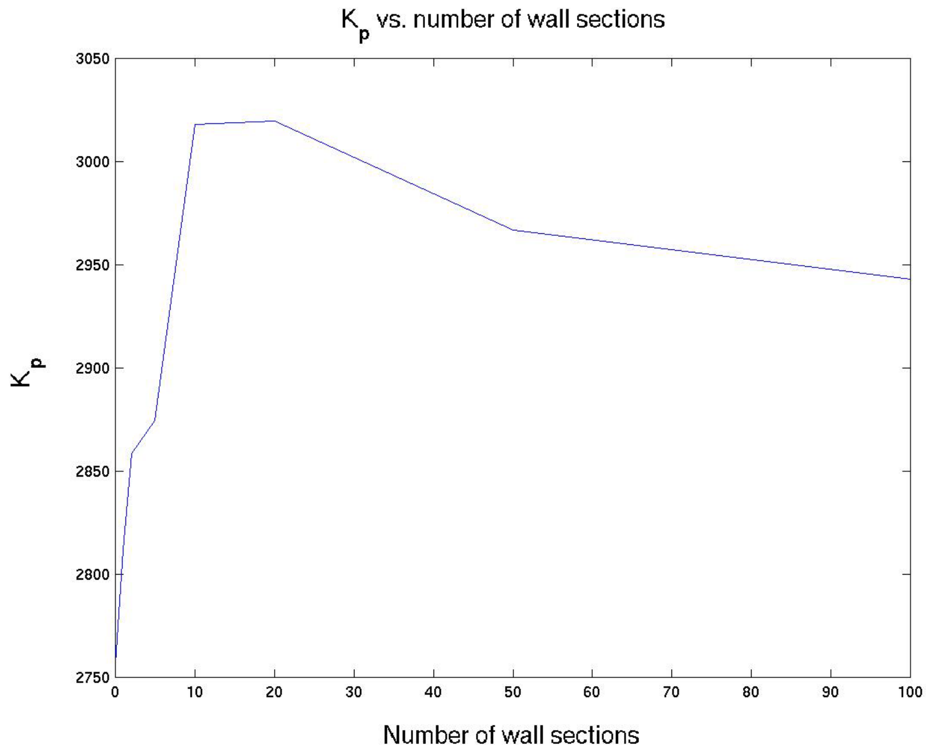

Figure 19 - K_p vs. number of wall sections

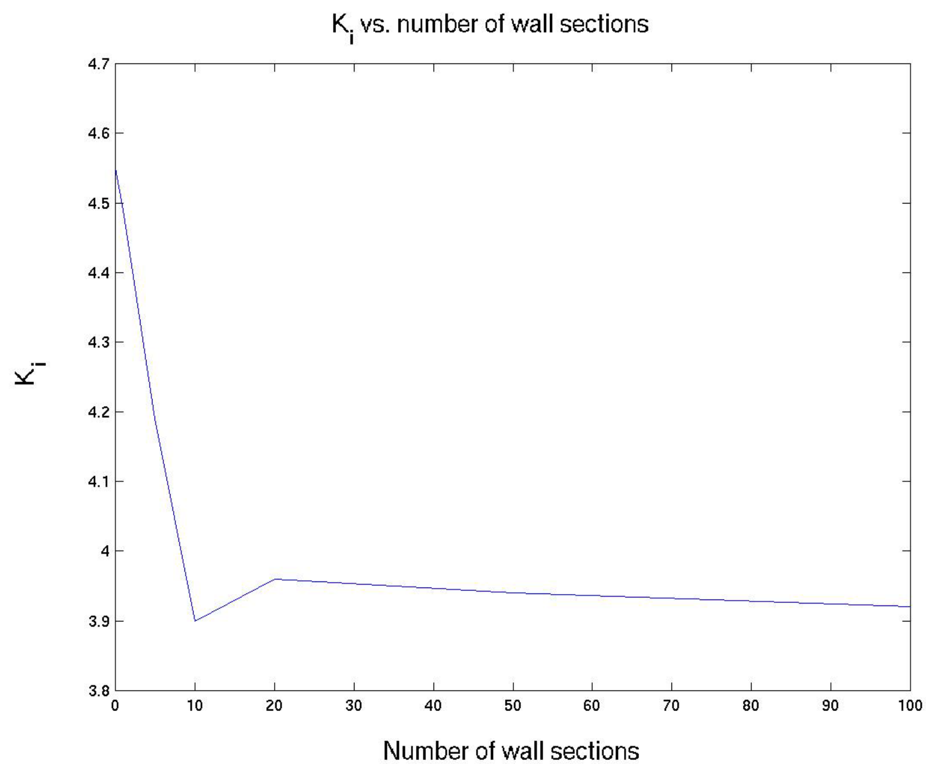

Figure 20 - K_i vs. number of wall sections

The PI gains found in Table 1 seem to approach finite limits with increasing finite wall elements, as illustrated in Figs. 19 & 20. This is explained by the fact that the finite difference approximation used in modeling the wall approaches Eqn. 27 as the number of wall elements increases. Therefore the limiting value for the PI gains should be used. However, the values for K_p and for K_i have fairly tight tolerances, with standard deviation of 96, mean ≈ 2900, and standard deviation 0.27, mean ≈ 4.2 respectively. As such, the PI gains may either be chosen according to Table 1 given the number of finite wall elements, or by the average value.

3.2.3 Industry Version for Performance Comparison

A simple PI controller was chosen for comparison to controllers with two gains used in industry. Therefore a purely digital, discrete-time two-term controller with the following control algorithm representing the industry controller must be derived:

$$u_{p_{k+1}} = \frac{100}{PB}\left[e_k + \frac{\tau}{T_i}\sum_{n=0}^{\infty} e_{k-n} \right]$$

where u_p = percentage of total control effort
PB = proportional band term
T_i = integral time term
τ = sampling interval
e = error (controller input)

Equation 46 - Industry Control Algorithm

In order to obtain a state space realization of the industry control algorithm, extensive algebraic manipulations are required. The first step in the derivation is to decrement the entire equation one step back in time:

$$u_{p_k} = \frac{100}{PB}\left[e_{k-1} + \frac{\tau}{T_i}\sum_{n=0}^{\infty} e_{k-n-1} \right]$$

Subtracting this equation from Eqn. 46 yields:

$$u_{p_{k+1}} = u_{p_k} - \frac{100}{PB} e_{k-1} + \frac{100}{PB} e_k + \frac{100}{PB}\frac{\tau}{T_i} e_k$$

Simplifying:

$$u_{p_{k+1}} = u_{p_k} + \frac{100}{PB}\left[\left(1 + \frac{\tau}{T_i}\right) e_k - e_{k-1} \right]$$

Equation 47 - Percent Control Effort Using Industry Control Algorithm

The control algorithm is in percent of full control effort. Therefore, in order to obtain an actual value in Watts or kW for the output of the controller, typical heat load values of a space are needed. A sampling of different building types yields an average maximum refrigeration cooling load value of 188.1 sq.ft./ton, or the equivalent metric value representing the maximum cooling intensity (MCI) of 201.2 W/m^2. The percentage of full control effort can be scaled by this value and multiplied by the square footage of floor space, A_{fl}, giving the actual output of the controller:

$$u_{k+1} = \frac{u_{p_{k+1}}}{100} MCI \cdot A_{fl}$$

Equation 48 - Actual Industry Control Effort

Combining Eqns. 47 & 48:

$$u_{k+1} = u_k + \frac{MCI \cdot A_{fl}}{PB}\left[\left(1+\frac{\tau}{T_i}\right)e_k - e_{k-1}\right]$$

Equation 49 - Final Industry Input/Ouput Controller Equation

The proportional band (PB) and integral time (T_i) terms govern the behavior of the controller.

Redefining Eqn. 49:

Let $K_p = \frac{MCI \cdot A_{fl}}{PB}$, and $K_i = 1 + \frac{\tau}{T_i}$, → $u_{k+1} = u_k + K_p[K_i e_k - e_{k-1}]$.

Equation 50 - Revised Industry Input/Output Controller Equation

Define the state-space representation with input e_{k-1} and output u_k: $\begin{bmatrix} x_{c_{k+1}} \\ u_k \end{bmatrix} = \begin{bmatrix} A_{c_d} & B_{c_d} \\ C_{c_d} & D_{c_d} \end{bmatrix}\begin{bmatrix} x_{c_k} \\ e_{k-1} \end{bmatrix}$

Because the controller is 1st order, an input/output equation version of the state-space representation can be obtained by using some algebraic manipulations:

$$u_{k+1} = A_{c_d} u_k + D_{c_d} e_k - (A_{c_d} D_{c_d} - C_{c_d} B_{c_d}) e_{k-1}$$

Matching with Eqn. 50, $u_{k+1} = u_k + K_p K_i e_k - K_p e_{k-1}$

Yields: $A_{c_d} = 1,\ D_{c_d} = K_p K_i,\ \text{and } A_{c_d} D_{c_d} - C_{c_d} B_{c_d} = K_p$

Free parameters exist, therefore let $C_{c_d} = K_p$, and $B_{c_d} = K_i - 1$, finally giving:

$$\begin{bmatrix} x_{c_{k+1}} \\ u_k \end{bmatrix} = \begin{bmatrix} 1 & K_i - 1 \\ K_p & K_p K_i \end{bmatrix}\begin{bmatrix} x_{c_k} \\ e_{k-1} \end{bmatrix}$$

However, because the input to the industry controller, e_{k-1}, is delayed one step, a redefined augmented state accounting for the delay is as follows:

$$\bar{x}_{c_k} = \begin{bmatrix} x_{c_k} \\ e_{k-1} \end{bmatrix}$$

With the two equations

$$x_{c_{k+1}} = x_{c_k} + (K_i - 1) e_{k-1}$$

$$e_k = e_k$$

An augmented two-state industry controller can be formed:

$$\bar{x}_{c_{k+1}} = \begin{bmatrix} x_{c_{k+1}} \\ e_k \end{bmatrix} = \begin{bmatrix} 1 & K_i - 1 \\ 0 & 0 \end{bmatrix}\begin{bmatrix} x_{c_k} \\ e_{k-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} e_k$$

$$u_k = \begin{bmatrix} K_p & K_p K_i \end{bmatrix}\begin{bmatrix} x_{c_k} \\ e_{k-1} \end{bmatrix}$$

Therefore, the final state equation representing the industry controller is:

$$\begin{bmatrix} \overline{x}_{c_{k+1}} \\ u_k \end{bmatrix} = \begin{bmatrix} 1 & K_i - 1 & 0 \\ 0 & 0 & 1 \\ K_p & K_p K_i & 0 \end{bmatrix} \begin{bmatrix} \overline{x}_{c_k} \\ e_k \end{bmatrix}$$

$$\text{where } K_p = \frac{MCI \cdot A_{fl}}{PB}$$

$$K_i = 1 + \frac{\tau}{T_i}$$

Equation 51 - Derivation of Industry Controller in State-Space Form

Applying Eqn. 41 and the same three step gain-tuning procedure outlined in Sec 3.2.2 can be used to form the closed-loop with the industry controller. The step response resulting from using these methods is as follows:

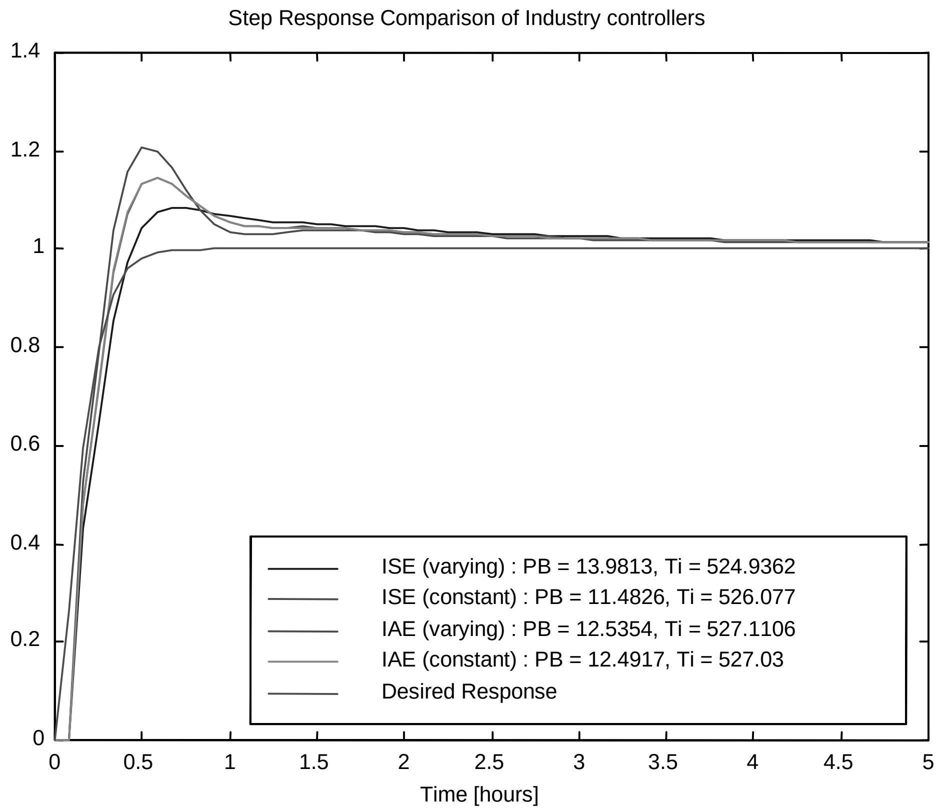

Figure 21 – Closed-Loop Step Response for Industry Controller

In Fig. 21, the industry controller shows a better IAE response than in Fig. 17 or 18 with the PI controller. Actual industry controller gains tuned using an IAE criterion method are available for comparison, but the values are much different than the gains arrived at above. This is due to the fact that the model used to tune the gains above was different than the space used to tune the gains provided in the following table (courtesy of Johnson Controls, Inc.):

VAV Zone PB & TI	%	Type
5P2F3ZON\VRH_SP01.AD_10	247.2	Clg Ti
5P2F3ZON\VRH_SP02.AD_9	6.6	Clg Pb
5P2F3ZON\VRH_SP02.AD_10	313	Clg Ti
5P2F3ZON\VRH_SP03.AD_9	4.7	Clg Pb
5P2F3ZON\VRH_SP03.AD_10	471.7	Clg Ti
5P2F3ZON\VRH_SP04.AD_9	9.8	Clg Pb
5P2F3ZON\VRH_SP04.AD_10	349.4	Clg Ti
5P2F3ZON\VRH_SP13.AD_9	7.1	Clg Pb
5P2F3ZON\VRH_SP13.AD_10	387	Clg Ti
5P2F3ZON\VCO_SP01.AD_9	1.5	Clg Pb
5P2F3ZON\VCO_SP01.AD_10	252.6	Clg Ti
5P2F3ZON\VCO_SP02.AD_9	2.6	Clg Pb
5P2F3ZON\VCO_SP02.AD_10	285.8	Clg Ti
5P2F3ZON\VCO_SP03.AD_9	4.5	Clg Pb
5P2F3ZON\VCO_SP03.AD_10	221.7	Clg Ti
5P2F3ZON\VCO_SP04.AD_9	3.4	Clg Pb
5P2F3ZON\VCO_SP04.AD_10	149	Clg Ti
5P2F3ZON\VCO_SP05.AD_9	15.5	Clg Pb
5P2F3ZON\VCO_SP05.AD_10	526.5	Clg Ti
5P2F3ZON\VCO-SP06.AD_9	13.6	Clg Pb
5P2F3ZON\VCO-SP06.AD_10	266	Clg Ti
5P2F3ZON\VRH_SP01.AD_11	2.3	BxHt Pb
5P2F3ZON\VRH_SP01.AD_12	1693.5	BxHt Ti
5P2F3ZON\VRH_SP02.AD_11	6.8	BxHt Pb
5P2F3ZON\VRH_SP02.AD_12	1104.4	BxHt Ti
5P2F3ZON\VRH_SP03.AD_11	1.1	BxHt Pb
5P2F3ZON\VRH_SP03.AD_12	246.7	BxHt Ti
5P2F3ZON\VRH_SP04.AD_11	11.9	BxHt Pb
5P2F3ZON\VRH_SP04.AD_12	596.5	BxHt Ti
5P2F3ZON\VRH_SP13.AD_11	7	BxHt Pb
5P2F3ZON\VRH_SP13.AD_12	337.9	BxHt Ti
5P2F3ZON\VRH_SP13.AD_13	7.1	SupHt Pb
5P2F3ZON\VRH_SP13.AD_14	228.2	SupHt Ti

Table 3 - Actual Industry-Tuned Gains

The first column indicates the type of area for which the gains were tuned. Zones 4 (SPO4) and 13 (SP13) are conference rooms. The other zones are open office spaces with short cubicles. The remaining columns indicate the Pb & Ti controller gain values and distinguish the type of cooling & heating used. From Fig. 21, it is evident that the ISE criterion method with varying trajectory reference exhibits the best step response characteristics. Therefore, the IAE criterion method can be to tune the industry gains successfully, but the ISE varying trajectory method is best for tuning both the PI controller & the industry controller.

3.2.4 Initial Conditions & Standard Simulation Parameters

The initial conditions for the simulation are obtained by finding the equilibrium state of the closed-loop system. The state space realization for the closed-loop from Eqn. 41 can be compactly represented with an augmented state defined as:

$$\tilde{\mathbf{x}}_k = \begin{bmatrix} \mathbf{x}_k \\ x_{c_k} \\ x_{lag_k} \end{bmatrix}, \text{ and } \tilde{\mathbf{u}}_k = \begin{bmatrix} \mathbf{w}_k \\ r_k \end{bmatrix}, \text{ therefore } \begin{bmatrix} \tilde{\mathbf{x}}_k \\ y_k \end{bmatrix} = \begin{bmatrix} \mathbf{A_{cl}} & \mathbf{B_{cl}} \\ \mathbf{C_{cl}} & \mathbf{D_{cl}} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{x}}_k \\ \tilde{\mathbf{u}}_k \end{bmatrix}$$

Equation 52 - Compact Closed-Loop Representation

In steady-state equilibrium, the discrete-time state will level off to some steady state, $\tilde{\mathbf{x}}_{ss} = \tilde{\mathbf{x}}_{k+1} = \tilde{\mathbf{x}}_k$. From Eqn. 52, $\tilde{\mathbf{x}}_{k+1} = \mathbf{A}_{cl}\tilde{\mathbf{x}}_k + \mathbf{B}_{cl}\tilde{\mathbf{u}}_k$, can be expanded: $\tilde{\mathbf{x}}_{k+1} = \mathbf{A}_{cl}\tilde{\mathbf{x}}_k + \mathbf{B}_{r_{cl}} r_k + \mathbf{B}_{w_{cl}} \mathbf{w}_k$. Hence, the equilibrium state can be determined, given $\mathbf{I} - \mathbf{A}_{cl}$ is invertible: $\tilde{\mathbf{x}}_{ss} = (\mathbf{I} - \mathbf{A}_{cl})^{-1}\mathbf{B}_{cl}\tilde{\mathbf{u}}_k$. Starting the simulation at equilibrium conditions prevents large initial transients from occurring and potentially biasing the complaint data logged during simulation. Therefore, the initial condition will take on the value of the equilibrium state, $\mathbf{x_0} = \tilde{\mathbf{x}}_{\mathbf{ss}}$. The initial state of the plant, controller, and sensor can be parsed out from $\mathbf{x_0}$, given the state definition in Eqn. 52.

The standard simulation parameters are as follows:

Parameter		**Value**
Sampling Interval		5 min
Number of Wall Sections		5
Control Strategy		PI
Controller Gains	K_p	2874.6058
	K_i	4.192
Desired Step Response Settling Time Used in Gain Tuning		20 min
Sensor Lag		7.5 min

Table 4 - Standard Simulation Parameters

The sampling interval is 5 minutes so that it will be small enough to prevent events from being missed. In addition, smaller sampling intervals are advantageous due to having the option of using different analytical methods resulting in input noise covariances that are within 0.1 % error of each other. The limit on the length of the sampling interval is constrained by the hourly outside weather temperature data. Therefore, sampling intervals of an hour or under can be used by linearly interpolating intermediate data points. From Table 2, it appears that the PI gains for 5 finite elements is closest to the average value for all of the PI gains logged. The PI control strategy is chosen in lieu of the industry controller because it is a lower order controller. The desired step response settling time used in the PI gain tuning procedure is 20 minutes to provide for a reasonably quick response time in heating the space temperature by 1 degree. Finally, the time lag of the thermostat sensor measuring the space temperature is 7 - 8 minutes, particularly for thermostats used in industry. Hence the simulation parameters are set accordingly.

3.3 Finite State Machines

Now that the foundation for building, hot & cold complaint temperatures has been laid out, the framework that determines how the thermostat setting is changed and data logging of the performance metrics occurs can be developed. The interactions among these temperatures are sometimes quite subtle, and there are many different elements that go into determining the manner in which the thermostat settings vary, ultimately affecting the resulting performance metrics. Therefore, the entire scheme should be set in place with a formal logic design, via finite state machines. This section covers the four basic finite state machines used to implement both new & industry thermostat setting policies.

3.3.1 Complaint Detection

The method of simulating complaints and recording associated performance metrics was introduced in Sec. 2.3.2. Recall how the two basic performance metrics were logged from up & downcrossings among the three processes illustrated in Fig. 5. This section (Sec 3.3.1) is relevant to new and industry strategies because the annual number of complaints performance metric is recorded for both. A state transition diagram illustrating the detection of complaints is as follows:

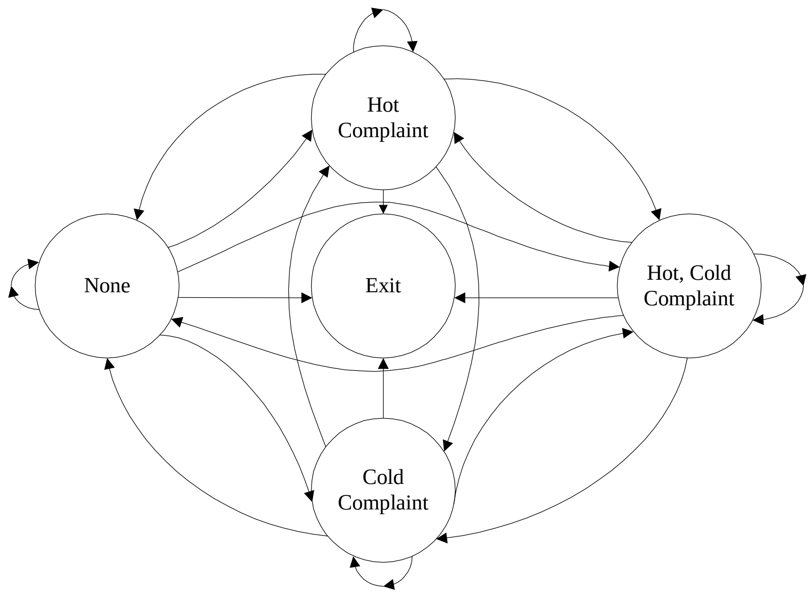

Figure 22 - Complaint Detection State Transition Diagram

Each circle in Fig. 22 represents a state, and the arrows illustrate different possible transitions. There are five states; the first is the none state, meaning that there are no complaints that have been detected. In simulation, the distinct hot and cold complaint states signify that a complaint condition exists, and the subsequent up or down crossing to end the complaint condition has not been detected yet. There is a state that attempts to capture the rare circumstance in which both hot & cold complaint conditions exist simultaneously, given by the hot & cold complaint state. Finally, the exit state denotes that the end of the simulation has been reached and appropriate simulation cleanup takes place.

The initial state upon entry into simulation is always the none state. This is guaranteed by polling for initial conditions until they satisfy the requirements of being in the none state. This means that the hot complaint temperature must be greater than the building temperature, and the building temperature must be greater than the cold complaint temperature. The primary indicator of crossings between building & complaint temperatures is the sign of the difference between current building & complaint temperatures. As an example, if in the none state, subtracting the building temperature from the cold complaint temperature should always yield a negative value. If this value changes sign while in the none state, a crossing has been detected and transition to another state should take place. For formality, each state has been broken into three distinct sections: the entry section, action section, and test/exit section. The entry section is executed once on entry to the state, the action section is executed on every scan of the state, and the test/exit section tests the condition for a transition and executes associated code if the test for the transition passes (Auslander, 1997, p. 14).

In the none state, the entry section contains code which prints out diagnostic data to standard output if the option is selected by the user, in addition to keeping track of the fact that the none state was the last state occupied for future time steps (hereto after referred to as state tracking). There are no commands that need to be executed on every scan of the state; therefore no action section is necessary. The test/exit section first determines if there is a pending simulation termination, in which case a transition to the exit state occurs on the subsequent time step. The remaining tests check for crossings between building & complaint temperatures that occurred between the previous time step and the current one. Using the sign differences between the current complaint & building temperatures as discussed earlier, it can be determined whether or not a transition must take place. Only the current complaint & building temperatures are used to make this determination, as opposed to temperatures at both the current & previous step. Extenuating circumstances may exist where the complaint & building temperatures are equal for any number of time steps. Because the number of time steps for which the complaint & building temperatures may be equal is unknown, using the temperatures at the previous step is not necessarily helpful in determining if a crossing has taken place. Rather, the knowledge of being in the current state in addition to any current equalities or inequalities of building & complaint temperatures is used to make this determination. As such, the remaining transition states are set according to the following table:

ΔT_c	ΔT_h	Transition State
>0	≥0	Cold Complaint
≤0	<0	Hot Complaint
>0	<0	Hot & Cold Complaint

Table 5 - Summary of Transitions from None State

Definitions: $\Delta T_c = T_c - T_r$, $\Delta T_h = T_h - T_r$

where T_c = Cold complaint temperature, T_h = Hot complaint temperature, T_r = Room Temperature

The detection of a crossing is based only upon strict equalities, and no transitions are made upon tight inequalities. Tight inequalities represent situations where there is no cause for a transition, or a transition might occur shortly due to the complaint & building temperatures being equal. An equality is not sufficient evidence of there being a crossing. Only subsequent steps resulting in non-zero sign changes of ΔT_c or ΔT_h determine if a transition should take place or not. In addition to state transitions, flags are set for communication with other finite state machines. These flags will activate timers associated with complaint recovery periods or the setback countdown timer for the new strategy. Linear interpolation between previous & current temperatures involved in the crossing and associated time steps is used to determine the actual start of a complaint recovery period in lieu of the closest time step. Finally, in the case that none of the above conditions for transitions are met and no simulation termination is pending, i.e. where $\Delta T_c \leq 0$ and $\Delta T_h \geq 0$, the state self-transitions, remaining in the none state.

In the cold complaint state, the entry section contains code whose execution is dependent upon the last state that was occupied. This illustrates the importance of state tracking, as performed in the none state. If the previous state was hot complaint or none, the thermostat setting is adjusted upwards using the current parameters of the appropriate strategy. The thermostat setting change will not occur in the same time interval that the complaint is detected, but takes one time step more. The setting change occurs in the subsequent step to it being detected so that the state transition can take place. As a result, the thermostat setting change will occur at most two time steps from the actual crossing because most crossings occur in between time steps, and it takes up to one time step for the crossing to be detected. The logic adds an implicit delay to adjusting the thermostat setting to a new value when receiving a complaint. However, in practice this is more than often the case because if facility operators implement the policy manually, the delay may even surpass the 5 - 10 extra minutes required to adjust the thermostat setting in simulation. Following the thermostat setting adjustment, a counter keeps track of how many cold complaints have taken place throughout the simulation, and optionally, diagnostic data prints to standard output. Additional diagnostic data is stored in an audit trail, particularly when events such as complaints occur to justify it. The audit trail logs the time in days, the status of all of the finite state machines, and all temperatures prior to and after the event occurred. In addition, a tracking variable stores the fact that the cold complaint state had been the last state occupied for future time steps. If the previous state was hot & cold complaint, the end of a hot complaint recovery period is implied. In this case no code other than diagnostic information to standard output and state tracking is required.

The action section of the cold complaint state contains code used for printing data to standard output. It will also shut off a datalogger whose function is to write to file a 1-year circular buffer of time-series building, complaint, and thermostat temperature data if a CRP exceeds a certain value. The reason for investigating long CRP's has to do with finding the complaint model breakdown threshold, which will be discussed in the conclusion. The test/exit section of the cold complaint state contains tests similar to the none state, for determination of pending state transitions. Again, the first test determines if there is a pending simulation termination, in which case a transition to the exit state occurs on the subsequent time step. The remaining tests check for crossings between building & complaint temperatures, and transition states are set according to the following table:

ΔT_c	ΔT_h	Transition State
<0	≥0	None
≥0	<0	Hot & Cold Complaint
<0	<0	Hot Complaint

Table 6 - Summary of Transitions from "Cold Complaint" State

Some of the actions that follow immediately after setting the new transition state are similar to the ones described in the none state. Flags are set in order to communicate with other finite state machines for activation or termination of timers associated with complaint recovery periods or the setback countdown timer. Linear interpolation involving both the previous & current temperatures also takes place to determine exact upcrossings & downcrossings that signify the end of the current cold CRP and/or the beginning of a new hot CRP. However, when transitioning to the none state or the hot complaint state, other actions take place as well because transitioning to these states from the current cold complaint state implies the end of the current cold CRP. Therefore, tracking of total simulation time in a cold complaint recovery period needs to take place to record the average CRP performance metric. Additionally, error messages related to the logging of time in the cold CRP are printed to standard output, and to file if the cold CRP exceeds a certain value. In the case that none of the above conditions for transitions are met and no simulation termination is pending, where $\Delta T_c \geq 0$ and $\Delta T_h \geq 0$, the state self-transitions, remaining in the cold complaint state. The hot complaint state code contains the same details as the cold complaint state, except that everything is reversed. The thermostat setting is adjusted downwards instead of upwards, appropriate variables are used in recording hot CRP's & detecting crossings, and the table used to determine transitions in the test/exit section of this state is as follows:

ΔT_c	ΔT_h	Transition State
>0	≤0	Hot & Cold Complaint
≤0	>0	None
>0	>0	Cold Complaint

Table 7 - Summary of Transitions from "Hot Complaint" State

In the case that none of the above conditions for transitions are met and no simulation termination is pending, where $\Delta T_c \leq 0$ and $\Delta T_h \leq 0$, the state self-transitions, remaining in the hot complaint state.

Similar to the previous states, the entry section of the hot & cold complaint state contains code whose execution is dependent upon the last state that was occupied. If the previous state was cold complaint, then the thermostat setting is adjusted downwards, using the current parameters of the appropriate strategy. Transitioning from the cold complaint state to the hot & cold complaint state implies that a hot complaint has occurred, hence the reason for the negative thermostat adjustment. Similar to previous states, a counter keeps track of how many hot complaints have taken place throughout the simulation, diagnostic data is optionally printed to standard output, and an audit trail prints the appropriate information to file. Also, state tracking for the hot & cold complaint state takes place. If the previous state was hot complaint, then the thermostat setting is adjusted upwards using the current parameters of the appropriate strategy, etc. If the previous state was none, then no thermostat setting adjustment takes place at all, because progression from the none state to the hot & cold complaint state implies that both hot and cold complaints were detected simultaneously. In this case, the thermostat setting is left as is, rather than adjusting it upwards or downwards, because doing so would respond to only one of the complaints. Finally, the counters keeping track of hot & cold complaints are incremented, diagnostic data again is optionally printed to standard output, audit trails for both hot & cold complaints are logged, and state tracking for the hot & cold complaint state is performed.

The action section of this state contains code used for printing to standard output, and shuts off the datalogger writing the time-series temperature data to file if either a hot or cold CRP exceeds a certain value. The test/exit section of the hot & cold complaint state contains similar tests as before, for determination of pending state transitions. Again, the first test determines if there is a pending simulation termination, in which case a transition to the exit state occurs on the subsequent time step. The remaining tests check for crossings between the building temperature & complaint temperatures that occurred between the previous time step and the current one. They are set according to the following table:

ΔT_c	ΔT_h	**Transition State**
<0	≤0	Hot Complaint
≥0	>0	Cold Complaint
<0	>0	None

Table 8 - Summary of Transitions from "Hot & Cold Complaint" State

In case none of the above conditions for transitions are met and no simulation termination is pending, where $\Delta T_c \geq 0$ and $\Delta T_h \leq 0$, the state self-transitions, remaining in the hot & cold complaint state. All of the actions that follow immediately after setting the new transition states are similar to the ones described before. Flags are set in order to communicate with other finite state machines for activation or termination of timers associated with complaint recovery periods and/or the setback countdown timer. Linear interpolation involving both the previous & current temperatures takes place to determine exact upcrossings & downcrossings signifying the end of either or both hot & cold CRP's. The total simulation time in a hot or cold complaint recovery period is tracked, and error messages related to logging CRP time are printed to standard output, and to file if the CRP exceeds a certain value. All of these actions are performed regardless of the pending state transition because from the hot & cold complaint state a CRP is always ending, whether it is a hot CRP, cold CRP, or both.

The entry section of the exit state is the only section with any written code. No action or test/exit section is required because this state is solely for cleanup. As before, the entry section of the exit state contains code whose execution is dependent upon the last state that was occupied. All previous states determine which timers & associated flags need to be turned off, and diagnostic data is printed to standard output when the option is selected to confirm that a clean exit was completed.

3.3.2 Complaint Recovery Periods

In order to distinguish between the detection of complaints and timers needed to keep track of CRP's, separate finite state machines to perform these functions individually for hot & cold CRP's must be created. This section is relevant to both new and industry strategies because the average complaint recovery period performance metric is logged in both modules. A state transition diagram to illustrate the different states & transitions required to keep track of both hot & cold CRP's and implement the two timers is as follows:

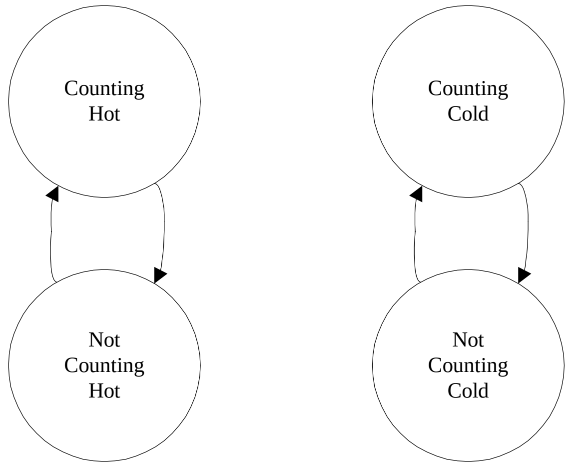

Figure 23 - Complaint Recovery Period State Transition Diagrams

There are two independent finite state machines shown in Fig. 23. As before, each circle represents a state that the finite state machine may occupy, and the arrows illustrate transitions between states. The counting states indicate that a complaint has been triggered or a complaint condition currently exists that has not yet been resolved. Hence, the time to resolve it is being “counted”. The not counting states indicate that no complaints have been detected, hence no counting is required to keep track of CRP lengths. The actual recording & logging of the actual CRP data occurs in the complaint detection finite state machine. These two complaint recovery period finite state machines are only a formality to ensure that important events are kept track of throughout the simulation.

Simulation never starts in a complaint condition, so the initial state upon onset is the not counting state. Similar to the complaint detection FSM (finite state machine), the entry section of the not counting state contains code whose execution is dependent upon the last state occupied. A CRP counter tallies the number of completed CRP’s if the counting state was previously occupied. Moving from the counting state to the not counting state implies that a CRP has ended. Optionally, diagnostic data can be printed to standard output, and an audit trail can write the specific details of the end of the complaint recovery period to file, and state tracking of the not counting state is performed. The only possible previous state is the counting state, therefore the only other possibility is that the not counting state is in self-transition mode. Because there is no code associated with the self-transition, there are no explicit arrows shown to indicate this in Fig. 23. The not counting state contains no action section, but the test/exit section does. In it the CRP start timer flag described in the complaint detection FSM resets, and transition to the counting state occurs in the subsequent time step. The entry section of the counting state performs the state tracking function. No action section exists, and the test/exit section serves only to reset the CRP end timer flags and transition to the not counting state.

3.3.3 Thermostat Setpoint Logic - Current Practice vs. New Strategy

Because the thermostat setting policies are implemented in finite state machine code, this section covers those two strategies in detail. As mentioned in Sec. 2, the optimized strategy is based on the magnitude & duration minimizing both performance metrics. An unbiased comparison of the metrics resulting from simulating current practice to the optimized strategy metrics then follows. In order to perform the optimization, a two-dimensional grid parametrized by magnitude & duration values needs to be established. These values are deterministic in nature, so each grid point will take on only the given value throughout the entire duration of that simulation. When the thermostat setting is changed, heuristics dictate that after a cold complaint is detected, the formula to compute the new thermostat setting is r = r + mag. This increases the thermostat setting in response to the cold complaint. After a hot complaint is detected, the formula to compute the new thermostat setting is r = r – mag, where r = the thermostat setting, and mag = the magnitude of the setpoint change. The new strategy's policy includes a setback portion as well, which uses a duration that counts down from some assigned value to zero. After the duration countdown comes to zero, the new strategy policy returns the thermostat setting to its original value. Hence, the formula for computing the new thermostat setting simply becomes $r = r_{orig}$. For comparison & analysis of the current practice policy with the new strategy, a one-dimensional grid for the industry strategy needs to be established. The industry policy is limited only to one parameter because there is no duration as in the new strategy. There is no forced setback of the thermostat setting after a given duration, because the new setting is changed for an indefinite amount of time until a complaint of the opposite type occurs. Remaining differences between the two thermostat setting policies are highlighted in this section by giving a detailed description of the statistical methods used in the industry strategy.

As mentioned in the introduction, the current practice thermostat setpoint policy is based on having no apriori knowledge of setpoint magnitude changes, so in simulation it is modeled as a random process. Setpoint changes on average tend to remain small in magnitude, with larger setpoint changes being less frequent. As a result, this information is used to select the appropriate probability distribution which most accurately represent the statistical model of the industry response to complaints, a one-sided Gaussian distribution. A standard Gaussian probability density function with mean = 5 °F and a variance of one is as follows:

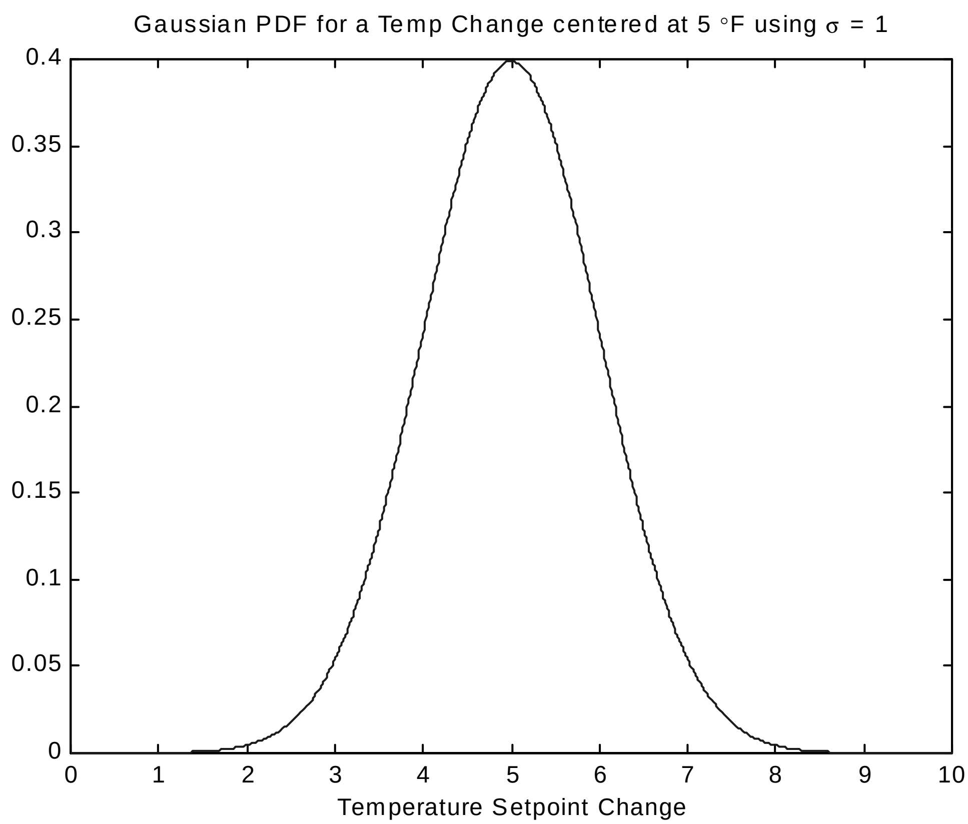

Figure 24 - Gaussian PDF, σ = 1

If the two-sided Gaussian distribution is selected as the governing PDF to represent the industry response, most of the setpoint changes selected randomly from the distribution would fall within a standard deviation on either side of the mean, between 4 and 6 °F. In summarizing the similarities between the two strategies from Sec. 2.2, it was mentioned that the direction of the setpoint change is always opposite to type of thermal sensation complaint. If a hot complaint is received, the appropriate response is to decrease the thermostat setting, and if a cold complaint is received the appropriate response is to increase the thermostat setting. However, if a Gaussian distribution such as the one shown in Fig. 24 is selected, there is a remote possibility that some temperature changes will be negative. This is contrary to the what the policy heuristics have outlined, for the direction of the setpoint change be opposite to the type of complaint detected. So by using the two-sided Gaussian PDF, negative setpoint changes may occur when they should be positive. To illustrate: if a hot complaint is generated, the formula used is r = r – sd, where the "sd" term must always be positive regardless of the statistics. If the "sd" term is negative because of the statistical distribution, the heuristics become biased due to cancellation of the minus sign in the formula, resulting in a positive increase. When hot complaints occur, the temperature change shouldn't be positive. Furthermore, this type of PDF is actually parametrized by two variables, not just one as should be the case.

The distribution governing the industry setpoint strategy can be slightly altered to be based on a different standard deviation, σ=3 but with same mean of 5 °F. Therefore it becomes more evident that inappropriate setpoint changes may occur by using the two-sided Gaussian PDF approach:

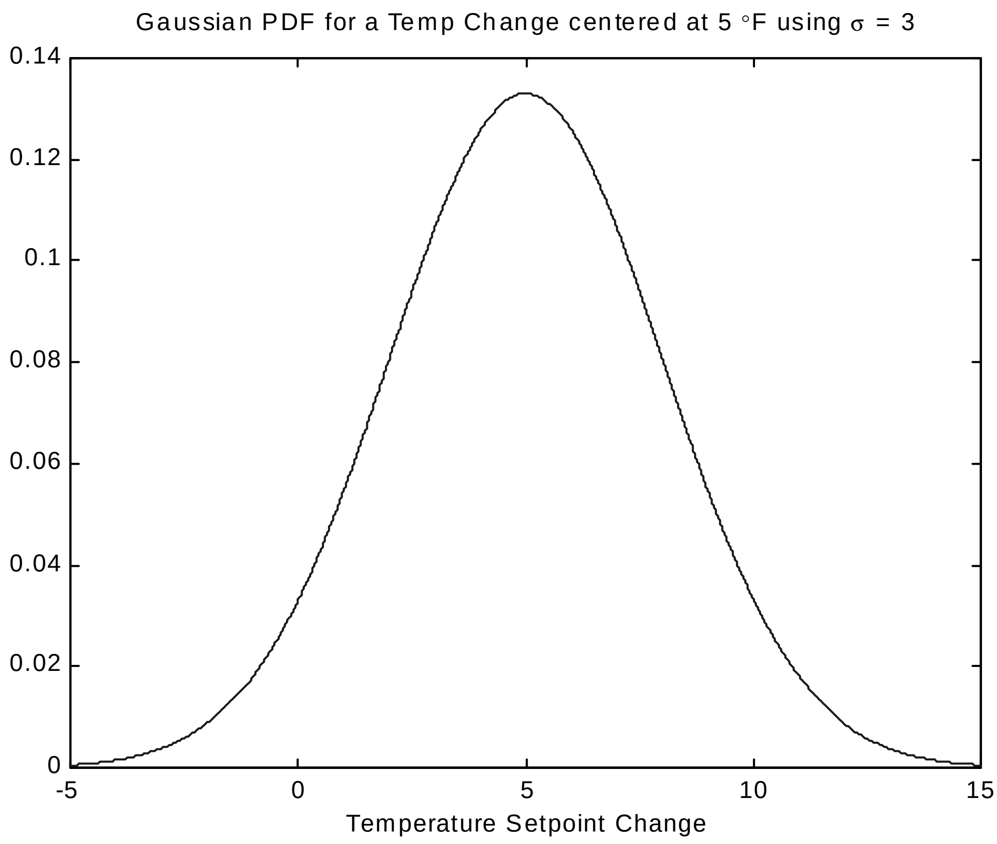

Figure 25 - Gaussian PDF, σ = 3

In Fig. 25, there is a reasonable probability that inappropriate temperature changes may occur when using a larger variance. Therefore it seems obvious that the solution to this problem is to use a one-sided Gaussian approach to prevent temperature changes of the wrong direction from occurring. A typical one-sided Gaussian PDF will have its own mean & variance, which can be computed analytically. However, because only a one-term representation of the industry strategy is desired, it's not important to have control over both the mean & variance of the one-sided PDF. Therefore the one-sided distribution can be treated as a standard Gaussian PDF with zero mean and variance set by the one-dimensional grid. The main difference is that the absolute value of the independent variable is used, giving the one-sided Gaussian PDF governed by a single parameter (shown with σ=5), as follows:

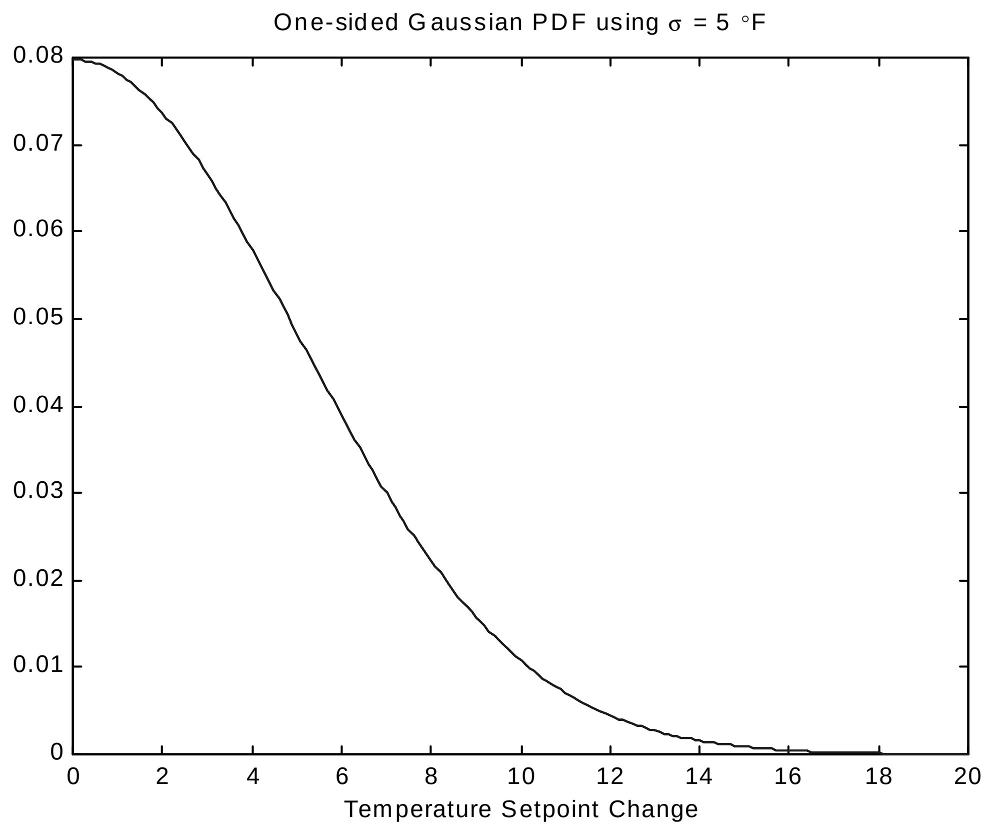

Figure 26 - One-sided Gaussian PDF with σ = 5 °F

This one-sided Gaussian PDF with σ = 5 °F has no directional problems, and represents one grid point for the current practice policy. This particular grid point would be compared to the new strategy that implements the thermostat setting formula $r = r \pm mag$, where mag = 5 °F. The new strategy is deterministic & throughout the entire simulation would have only magnitude changes of 5 °F for that particular grid point. However, the industry strategy is statistical in nature, & according to Fig. 26, the thermostat setting change for using σ = 5 °F may take on values ranging from 0 °F with high probability to values of 20 °F with low probability. Even so, it still is the best distribution to represent industry response to hot & cold complaints. Alternatives are other one-sided distributions such as the Poisson distribution, but these are not as easily realized in code. The reason why this distribution represents industry response so well is due to research in industry trends indicating the majority of setpoint changes having small magnitude, with larger setpoint changes being less frequent (Smothers, 1999). Therefore, the formula to be used in computing the industry thermostat setpoint policy is: $r = r \pm |N(0,\sigma^2)|$, where r = the thermostat setting, and $|N(0,\sigma^2)|$ = the absolute value of the number returned by a zero mean white Gaussian random variable with variance σ^2, (standard deviation = σ). To realize this formula in code, Eqn. 2 is used, giving the final form of the formula: $r = r \pm \sigma|N(0,1)|$. In summary, deterministic values of "mag" will be compared directly to the same statistical values of "σ". The formulas governing the two policies for the new & industry strategy are given by: $r = r \pm mag$ and $r = r \pm \sigma|N(0,1)|$, respectively.

3.3.4 Setback Counter

The setback counter FSM, unlike the others applies only to the new thermostat setting policy. The previous section reviewed the heuristics involved in new strategy, outlining how the duration counts down from an assigned value to zero. After the duration countdown comes to zero the new strategy dictates that the thermostat setting should be returned to its original value. There is an independent finite state machine formalizing this countdown sequence, and its state transition diagram is illustrated as follows:

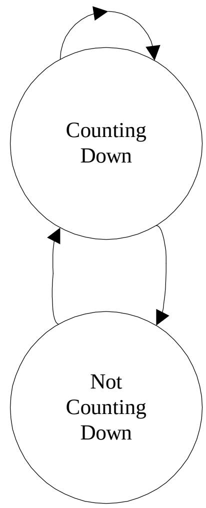

Figure 27 - Setback Counter State Transition Diagram

The counting down state shown in Fig. 27 indicates that a complaint has been triggered, so the countdown timer decrements to zero. The not counting down state indicates that no complaints have been detected, hence no countdown is required to keep track of the thermostat setback time. This is also the initial state because simulation never starts in a complaint condition. In the entry section of the not counting down state, only state tracking takes place. There is no action section associated with the not counting down state because there is no code that needs to be executed on every scan. The test/exit section has a single test that checks if the countdown timer flag has been set by the complaint detection FSM. If so, the flag is reset for future occurrences, the state transition to the counting down state is initiated, and the countdown timer variable is assigned the appropriate value. Because decrementing can only be performed with integers, the grid point value is approximated to the nearest integer in order to allow for a proper countdown sequence. This approximation is not a major flaw because if the policy is implemented without the aid of DDC, in all likelihood facility operators will not manually return the thermostat setting to its original value precisely at the time dictated by the policy.

Unlike other FSM's the entry section of the counting down state contains code whose execution is dependent upon a start setback timer flag, in addition to previous states occupied. If the not counting down state was occupied previously, standard state tracking is performed, and diagnostic data is optionally printed to standard output. Otherwise, the FSM is in self-transition mode indicated by the arrow in Fig. 27. If the start setback timer flag is set in this mode, a crossing has been detected prior to the end of the countdown, otherwise known as a countdown interruption. In this case, the thermostat policy dictates that the timer be reset and a new countdown take place. The complaint detection FSM will adjust the thermostat setting accordingly in response to the crossing detection, even though the setback has not occurred yet. This implies that immediate action of adjusting the thermostat in response to the pending complaint detection takes precedence over setting it back after countdown. In this case, if a large number of crossings occur in a very short period of time, depending upon which grid point is being simulated, the thermostat setting can be adjusted to an arbitrarily high or low setting. However, the very last complaint in the series of complaints will still adhere to the setback policy, and the thermostat setting will be adjusted back to the original value, not the value that it was at previously.

The entry section of the counting down state contains code which resets the start countdown flag, assigns the countdown timer variable its appropriate value, sets the start setback timer flag, and optionally prints diagnostic data to standard output only if in self-transition mode. Also, a flag that is used specifically to indicate countdown interruption is set. This flag is used to prevent restarting the countdown in the same time interval. For normal uninterrupted countdowns, initiation of the countdown is delayed a step and does not take place in the same time interval as crossing detection. Using this countdown interruption flag inserts the delay artificially. The action and test/exit sections are combined for the counting down state. If the countdown has not reached zero and the countdown interruption flag has not been set, the countdown timer decrements and appropriate diagnostic data is optionally printed to standard output. Otherwise, transition to the not counting down state takes place, and the new strategy kicks in, readjusting the thermostat setting to the original value, and the setback counter is incremented. Also, appropriate diagnostic data is optionally printed to standard output, in addition to an audit trail logging the specific details of the setback. Finally, the countdown interruption flag is reset, so that it can be used for future occurrences.

There is one final point to cover before ending this section, having to do with how the nominal value for the thermostat setting is arrived at. The nominal value is 72.56 °F, due to the fact that it is the temperature at which the simulation yields minimal performance metrics. The same simulation code used for the new & industry strategies was the basis for determining this nominal value, except that the grid search is over different static thermostat temperature settings, instead of varying thermostat setting policy parameters. The profiles of the performance metrics vs. the thermostat settings are illustrated in the following figure, to show how the optimal value for the thermostat setting was arrived at:

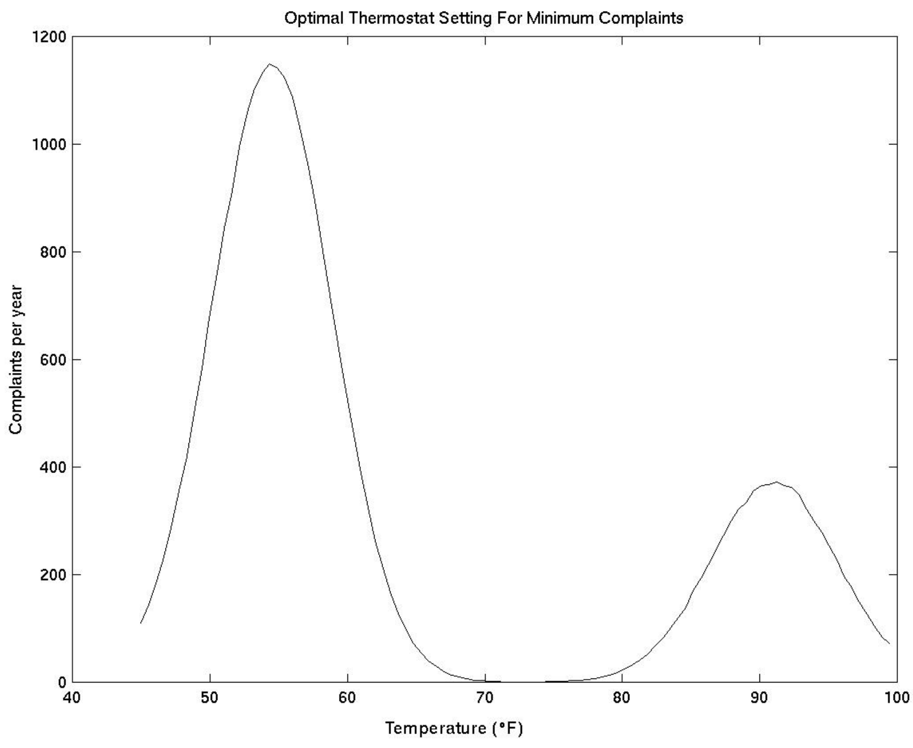

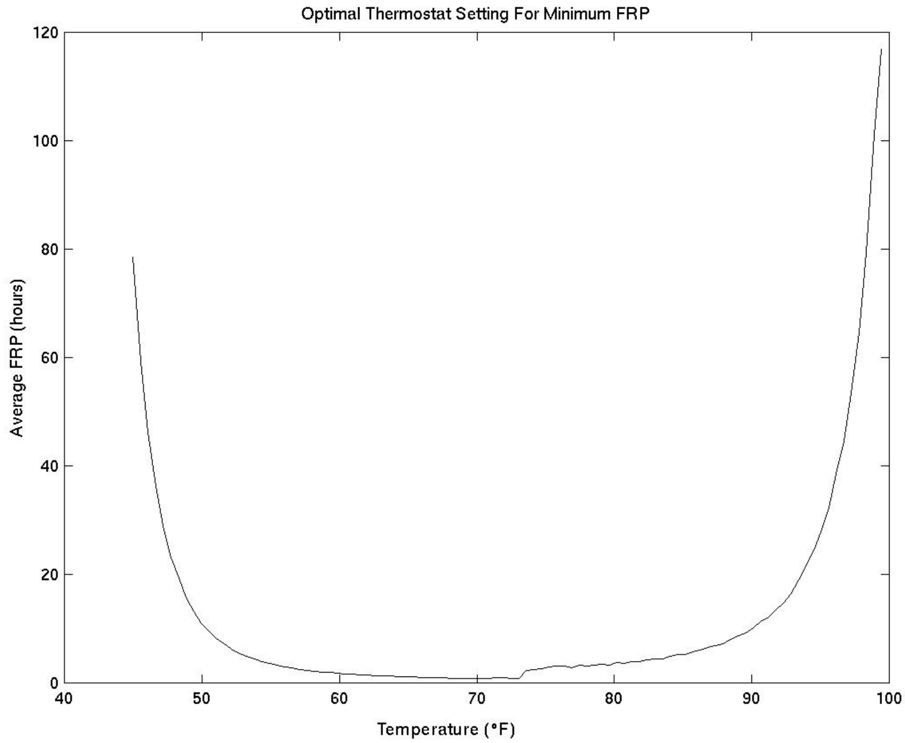

Figure 28 - Performance Metrics vs. Thermostat Setting

As seen in Fig. 28, both performance metrics exhibit convex behavior over the thermostat settings, making it easy to find the temperature at which the minimum values occur. In order to compute the optimal temperature, a percentage weighting on each metric is used. For the annual complaint metric shown in the top graph of Fig. 28, the temperature at which the minimum occurs is 73.05 °F, and on the bottom the temperature at which the minimum FRP occurs is 70.85 °F. Using a specific weighting between the performance metrics based 77.73 % on cost and 22.27% on response results in an optimal value of 72.56 °F. These temperatures are not perfectly accurate because the gridding of the temperature range might have been too coarse to obtain the temperature at which the exact minimums occurred. The gridding used was mainly to demonstrate the characteristic profiles of both performance metrics for a static thermostat setting over a large temperature range.

The reason for the shape of the first plot in Fig. 28 is due to the fact that the mean of the hot complaint temperature is 91 °F, and the mean of the cold complaint temperature is 54.5 °F. This accounts for the drastic peaks that are seen around these particular values. Because the thermostat setting is static, very little fluctuation other than the ambient disturbances affect tracking of the setpoint. Therefore, as the building temperature tracks these settings, it runs right through the most active sections of the hot & cold complaint levels. The peak for cold complaints is higher than the one for hot complaints because the rate of change of the cold complaint temperature is larger, hence there will be more cold complaints generated in the same time frame. There are no peaks in the average CRP metric plot, only a convex minimum. This is due to the fact that CRP's grow with increased building temperature difference away from nominal. When the thermostat settings are above or below the mean hot or cold complaint levels respectively, the crossings that determine when a complaint condition exists are very far removed from the crossings that determine where they end, resulting in longer CRP's.

The optimal thermostat setting of 72.56 °F will be the original setting used by both thermostat setting policies. Therefore, there is no bias in comparing the results of the two policies since the same initial thermostat setting is used to start with. However, the new strategy will always return to the nominal setting that has been optimized to yield minimum performance metrics, which may be one of the advantages of using this strategy.

3.4 Performance Metrics

Prior to comparing the two thermostat setting policies, the framework for presenting the results must be legitimized. The desired results are the performance metrics recorded in simulation, and are logged as follows: the annual number of complaints is computed by scaling the total number of complaints logged per grid point by the total simulation run time per grid point. The average CRP per grid point is computed by scaling the total time in both a hot & cold CRP status by the number of complaints detected for that grid point. This only poses a problem if there are no complaint logged for a particular grid point. Due to the stochastic nature of the simulations, uncertainty in the resulting performance metrics must be accounted for. Running the simulation once through the range of thermostat setting parameters in a grid search for an arbitrary length of time is not sufficient. The resulting performance metrics logged will be different for each different grid search. Although the results will be different, they also will be within a certain statistical tolerance of each other. These performance metrics will need to be optimized over in development of the new strategy, so it is very important that there is minimal error associated with them. Similarly, there will have to be minimum noise in the performance metrics resulting from running the industry strategy, in order for a comparison to be performed on unbiased grounds. There are many methods available for use in mitigating the stochastic uncertainty present in the final results. However, the final method used to reduce statistical uncertainty of the resulting performance metrics is presented as a viable alternative to these conventional methods.

3.4.1 Noise Minimization

In order to reduce stochastic uncertainty associated with the performance metrics, conditions for ergodicity must be met. The random process must be stationary in the strict sense, meaning that the probability density function is time-invariant. Alternatively, the system may have achieved stationarity in the wide/weak sense, meaning that the autocovariance function depends only on time differences in the experiment, and the expected value, or mean of the independent random variable in question, $m_x = E\{x(k)\}$ is time invariant. This mean or expected value can be represented as the ensemble average, computed by taking the mean of the random variable across several different samples of the same process at the same time running with distinct seeds so that the random sequences differ. The second condition dictates that the computed value for this ensemble average must agree with the time-averaged value, given by the equation (Tomizuka, ME233, 1999):

$$\overline{x}(k) \equiv \lim_{N \to \infty} \left(\frac{1}{2N+1} \sum_{i=-N}^{N} x(i) \right)$$

Equation 53 – Time average of random process

In Eqn. 53, x represents the independent random variable, and $\overline{x}(k)$ represents the running time-averaged value of the independent random variable over the duration of the experiment, N. The independent random variable x represents the number of complaints detected throughout the simulation. This definition is made for convenience because of the direct relationship to the annual number of complaints. It is an adequate representation for the entire model, without having to represent the second CRP performance metric as a random variable as well.

In order to ensure that the system is ergodic, an ensemble average as well as a time-averaged value must be computed during simulation and agree. The running average number of complaints detected over the duration of a simulation can be obtained by applying Eqn. 53. Establishing when this running time averaged value agrees with the ensemble average in simulation can be attempted, in trying to adhere strictly to the theoretical definition of ergodicity. The procedure to do so would first involve running the entire grid of thermostat setting parameters a fixed number of times. This number represents the total number of "grid simulations" in the ensemble to compute the ensemble average for the performance metrics at each grid point. Throughout the entire duration of each grid simulation, and in the midst of computing the ensemble averages of the performance metrics, a running time average using Eqn. 53 would be taking place to see when the ensemble average agrees with the time average for each grid point. Additionally, the random process must be stationary in the strict sense, in which case the probability density function must be time-invariant. However, because the true analytical nature of building occupant complaint behavior is not known well enough to be characterized with a probability density function, time-invariance of such a function can't be determined. Therefore a test for stationarity in the wide/weak sense stationarity in lieu of the strict sense must be used. The two tests for determining stationarity in the wide/weak sense are that the autocovariance function depends only on time differences in the experiment, and the expected value, or mean of the independent random variable in question, $m_x = E\{x(k)\}$ is time invariant. Once the mean has reached steady state, it an be considered time-invariant from that point forward. However, the second test for determining stationarity in the wide/weak sense involving the autocovariance function must still be applied. Computation of the autocovariance function and time dependencies is slightly more complicated. Therefore although theoretically sound, this procedure tends be quite painstaking. Furthermore, exact agreement between the ensemble & time averages, computation of the autocovariance function, and implementation of this method in code would be quite difficult. As such, this method is not adequately suited to making a determination of ergodicity. There are some analytical alternatives to this method, as well as some cookbook routines for general reduction of variance, particularly several Monte Carlo methods (Press, 1992, pp. 308, 316, 319) that aren't as exhaustive as this cumbersome ad-hoc method described. However, in lieu of trying any of these methods a viable alternative is presented in the next section.

3.4.2 Alternative to conventional methods for reduction of variance

As an alternative to the methods described in the previous section, a new approach to reducing stochastic uncertainty in performance metrics can be developed. It is based on a cost function defined as the ratio of the standard deviation to the mean of the random process in question. The independent variable associated with this cost function represents the number of complaints detected in simulation. A standard optimization problem therefore presents itself, given by the following equation:

$$\text{Cost} = \frac{\sigma[x]}{E[x]}\text{, and the optimization to be performed is: } \min_{x} \frac{\sigma[x]}{E[x]}$$

Equation 54 - Cost Function Describing Stochastic Uncertainty & Optimization

E[x] is the mean or expected value of the independent variable, and σ[x] is the standard deviation of x. Defining and minimizing the cost as the ratio of the two is certainly an adequate method to reduce the statistical uncertainty associated with the performance metrics, due to the relatively smaller value of the standard deviation with respect to the mean. The random variable describing the number of complaints detected throughout the simulation must be characterized with the most appropriate probability density function. In order to perform this optimization the random variable must be assigned a limiting distribution even though the true analytical nature of building occupant complaint behavior is not known well enough to do so. Nevertheless, the Poisson distribution seems to be the best choice due to the fact that it is meant to describe random processes dealing with events occurring during a specific time or in a specific area. The only possible issue with using this distribution is that there must exist next to no probability that two events or incidents could happen at the same time or place, and that they happen independently. Although the finite state machines allow for multiple complaints to occur simultaneously, there is very small likelihood of such a dual event happening, and the events are always independent. Some examples of random processes fitting into the description of a Poisson distribution are: accidents on a stretch of road, phone calls received in one day, bugs in computer code, number of light bulbs that will burn out on campus today, typos in textbooks, etc. Therefore, complaint detection during simulation seems to fit rather neatly into this classification of Poisson events.

Some basic facts about the Poisson distribution are as follows: α = average number of events or incidents per unit time or length, t = length of time or size of area, and Poisson parameter $\lambda = \alpha t$. The Poisson parameter, λ, is equal to the number of events, and also equal to both the mean, E[x], and the variance, σ^2. Therefore, the standard deviation is given by $\sigma = \sqrt{\lambda}$. The probability density function is formally given by:

$$p(x) = \frac{e^{-\lambda}\lambda^x}{x!}$$

Equation 55 - Probability density function for Poisson distribution

The cost function resulting from using the Poisson parameter, λ, can be computed as follows:
Cost = $\frac{\sigma[x]}{E[x]} = \frac{\sqrt{\lambda}}{\lambda} = \frac{1}{\sqrt{\lambda}}$. The optimization problem posed in Eqn. 54 can most likely be cast into some neatly derived closed-form analytical solution. However, in lieu of this method, there is a much more direct way to approach minimization of this cost, by simply setting a percentage relative tolerance to be the limiting factor in deciding when the cost is small enough, so that the simulation terminates. This percentage tolerance represents how much noise is permitted in the final results. Therefore, the procedure used in determining the final metrics is directly related to the number of complaints generated during simulation, λ. The cost function implies that the greatest reduction of statistical uncertainty is realized by the largest number of complaints being generated during simulation, meaning that simulations will have to run for a very long time.

The procedure for ensuring that the performance metrics logged for all grid points have statistical uncertainty below some pre-specified tolerance is to use the noisiest grid point in the entire grid as the worse case to base the cost function on. The following definitions are used for clarification:

$\hat{R} = \frac{1}{\sqrt{\hat{\lambda}}}$ = Cost for a single grid point

$\hat{R}_{wc} = \frac{1}{\sqrt{\hat{\lambda}_{wc}}}$ = Worst case cost (noisiest grid point)

R_d = Desired cost (tolerance to meet for termination of simulations)

The goal of minimization can be approximated by the following tight inequality: $\hat{R}_{wc} \le R_d$, or, the cost of the noisiest grid point must be less than or equal to the desired cost. From the basic Poisson relationship: $\hat{\lambda}_{wc} = \hat{\alpha}_{wc}\hat{L}$, where $\hat{\alpha}_{wc}$ = the complaint rate per simulation time length for the noisiest grid point, and $\hat{L}$ = the simulation time necessary to get the cost of the noisiest grid point under the desired cost. With some elementary algebra using this equation, the inequality, and the equation for $\hat{R}_{wc}$, a formula for $\hat{L}$ can be derived: $\hat{L} \ge \frac{1}{R_d^2\hat{\alpha}_{wc}} \Rightarrow \hat{L}_{min} = \frac{1}{R_d^2\hat{\alpha}_{wc}}$. $\hat{L}_{min}$ is the minimum simulation time required to get the cost of the noisiest grid point under the desired cost, which is not known prior to simulation. Initially, the grid search will start with some arbitrary simulation run length, L_0. However, this is almost always not enough time to get the cost under the desired level. In that case, additional time required to achieve the desired cost must be computed. If ΔL = the additional time required, and L_0 = the initial simulation time, then the total time required to get the cost under the desired level is $\hat{L}_{min} = L_0 + \Delta L$. $\hat{L}_{min}$ can be rewritten solely in terms of simulation parameters:

$$\hat{L}_{min} = \frac{1}{R_d^2 \hat{\alpha}_{wc}}$$

Equation 56 - Minimum total simulation time

The complaint rate, α, will not be the same for each grid point due to the fact that different thermostat setting policies are being tested. However, the complaint rate for the noisiest grid point from the initial grid run, $\hat{\alpha}_0$, will be the same as the complaint rate **for the same grid point** in subsequent grid runs performed for the additional simulation time required to achieve the desired cost. Yet there is no guarantee that the grid point exhibiting the most noise in the initial run, $\hat{\alpha}_0$ will be **the same grid point** that exhibits **the most noise** in subsequent grid runs. However, an approximation needs to made based upon the fact that the swing in complaint rates throughout the noisiest portion of the grid is not large enough to make a significant difference, such that $\hat{\alpha}_0 \approx \hat{\alpha}_{wc}$. Therefore, the basic Poisson relationship for the worst grid point in the initial run can be used: $\hat{\lambda}_0 = \hat{\alpha}_0 L_0$. With all of the above facts and the cost of the noisiest grid point of the initial grid run, given by:

$$\hat{R}_0 = \frac{1}{\sqrt{\hat{\lambda}_0}}$$

Equation 57 - Cost for noisiest grid point in initial grid run

A revised version of Eqn. 56 can be derived as follows:

$$\hat{L}_{min} \cong L_0 \left(\frac{\hat{R}_o}{R_d} \right)^2$$

Equation 58 - Revised Minimum Simulation Time

As captioned, this gives the total minimum simulation time required to get the cost below the tolerance. It is also strictly in terms of parameters that are immediately available during simulation. The extra time needed to get the cost below the tolerance can now be computed:

$$\Delta L = \hat{L}_{min} - L_0 \cong L_0 \left[\left(\frac{\hat{R}_0}{R_d} \right)^2 - 1 \right]$$

Equation 59 - Additional simulation time required

In simulation, it is easy to search through the grid and find the noisiest grid point on the initial grid run, giving the value of $\hat{R}_0$. The value of L_0 is immediately available since it's the initial simulation runtime length, and R_d acts as a simulation parameter specified by the user. Therefore, grid runs proceed repeatedly with additional simulation times dictated by Eqn. 59 until the noisiest grid point is below the specified tolerance. Often times not enough complaints are generated in the additional time computed by Eqn. 59, and several more grid runs need to be performed. In this case, the additional time required becomes smaller and smaller, making it more difficult for the correct number of complaints to be received in order to terminate the simulations. If this happens, a workaround has been implemented to set a minimum time for ΔL. If the computed value for ΔL exceeds this minimum time, then additional grid runs may proceed as necessary. However, if ΔL is below the minimum time, then in addition to using the computed value in Eqn. 59, a fixed amount of time known to be enough to generate a reasonable amount of complaints is added on to prevent infinite loops. The minimum time for ΔL used is 1 year, and the fixed amount of time known to be enough to generate a reasonable number of complaints is 10 years. Both of these values were based both upon apriori empirical & analytical results in addition to the fact that the sampling interval being used in all simulations is 5 minutes. The standard percentage tolerance used as the desired cost is 5% for the results presented in the next section. Also due to the fact that complaints are sometimes spaced far apart, total simulation time can sometimes run as long as 1500 years of computing time. However, this is as a direct result of the percentage tolerance chosen.

4 Results

4.1 Model Breakdown Threshold & Limitations

Due to the nature of the methods used to reduce statistical uncertainty, simulations are required to run for an extremely long time. As a result, many of the very rare circumstances mentioned earlier might indeed occur. One example not mentioned yet happens only when running the current practice model, or the new strategy model with very sluggish poorly tuned PI gains. The problem occurs when the thermostat setting logic dictates a very large setpoint change, or a number of moderate setpoint changes in a short period of time. Therefore, any subsequent change not large enough may result in the building temperature getting stuck above or below the hot or cold complaint levels, respectively. An illustration of this specific type of behavior for the current practice model is shown in Fig. 29:

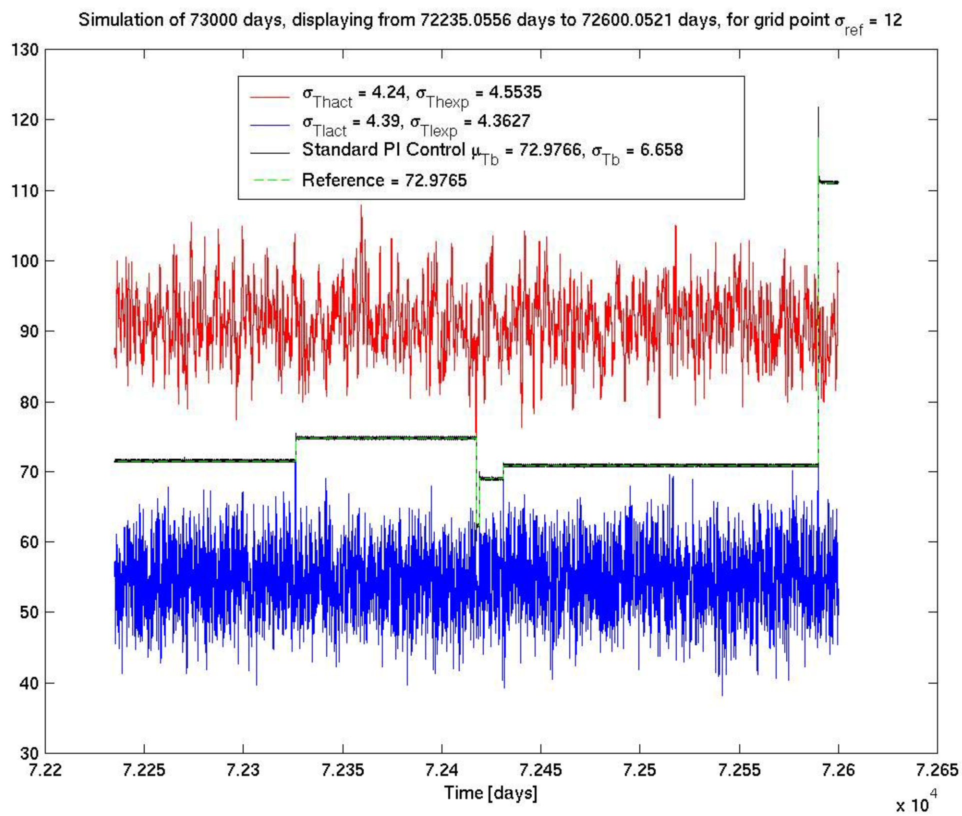

Figure 29 - Example of an Aberration

In Fig. 29, the building temperature is shown in black, hot & cold complaint temperatures are shown in red & blue, respectively, and the reference thermostat setpoint is given by the green dotted line.

Sample statistics for these processes are also shown on the legend. The subscript 'T' refers to temperature, 'l' refers to low or cold complaint level, and 'h' refers to high or hot complaint level. The subscript 'b' refers to building, 'act' denotes actual or empirical data, and 'exp' refers to experimental or sample data run in the simulation. Combinations of these subscripts explain each of the statistics provided in the legend. The building temperature follows the thermostat setting according to the current practice model. However, due to the fact that a large setpoint change is followed by a smaller one for $\sigma = 12°F$, the building temperature appears to get stuck above the hot complaint level towards the end of the years' worth of time-series data stored in the circular buffer shown. Zooming in on the area where this aberration occurs, the sequence of events causing this situation can be shown more clearly:

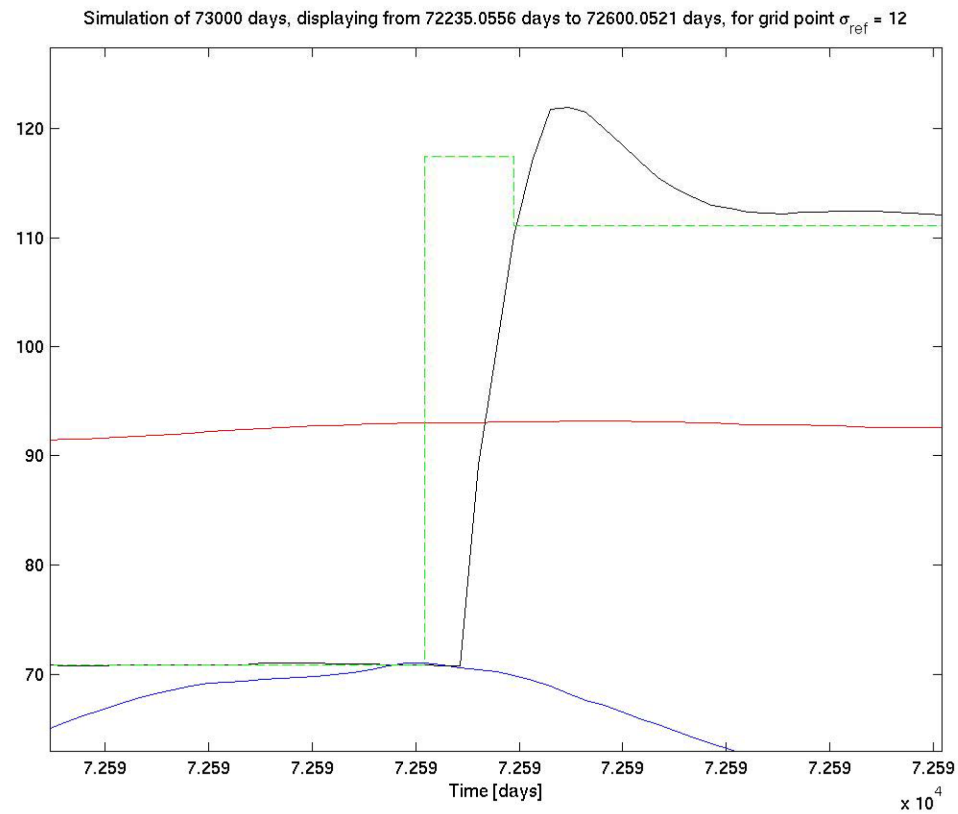

Figure 30 - Blow up of Sample Aberration

This type of behavior is indicated by long complaint recovery periods, due to the fact that the building temperature rises above the hot complaint temperature to meet the new setpoint. As is does, the complaint recovery period timer is started by the finite state machine shortly after the hot complaint is triggered. The downcrossing of the building temperature with the hot complaint level will not happen until the statistics of the hot complaint temperature cause a spike upwards to generate the end of the CRP. This event may not happen for an extremely long time, and all the while the CRP will continue to be timed. Therefore, in order to obtain a sample snapshot of the time series data containing the problem area, CRP's over a pre-specified large value are checked for continuously. Logging of this data is stored in order to determine at which magnitude, or σ value the anomalies first start to appear. The pre-specified CRP value is chosen heuristically as 24 hours, with no analytical basis. Using heuristics for the limiting CRP value is reasonable because a complaint condition should never exceed a day, in which case the building temperature would be stuck above or below the hot or cold complaint levels, respectively. In the specific case of the previous two figures, Fig. 29 & 30, the threshold CRP value used was 10 days as opposed to 24 hours, so that the aberration of the building temperature getting stuck would be clearly visible. Fig. 31

depicts the number of current practice simulation runs out of 12 on an interval of grid points from 0 to 12 °F for which the aberration described above occurs.

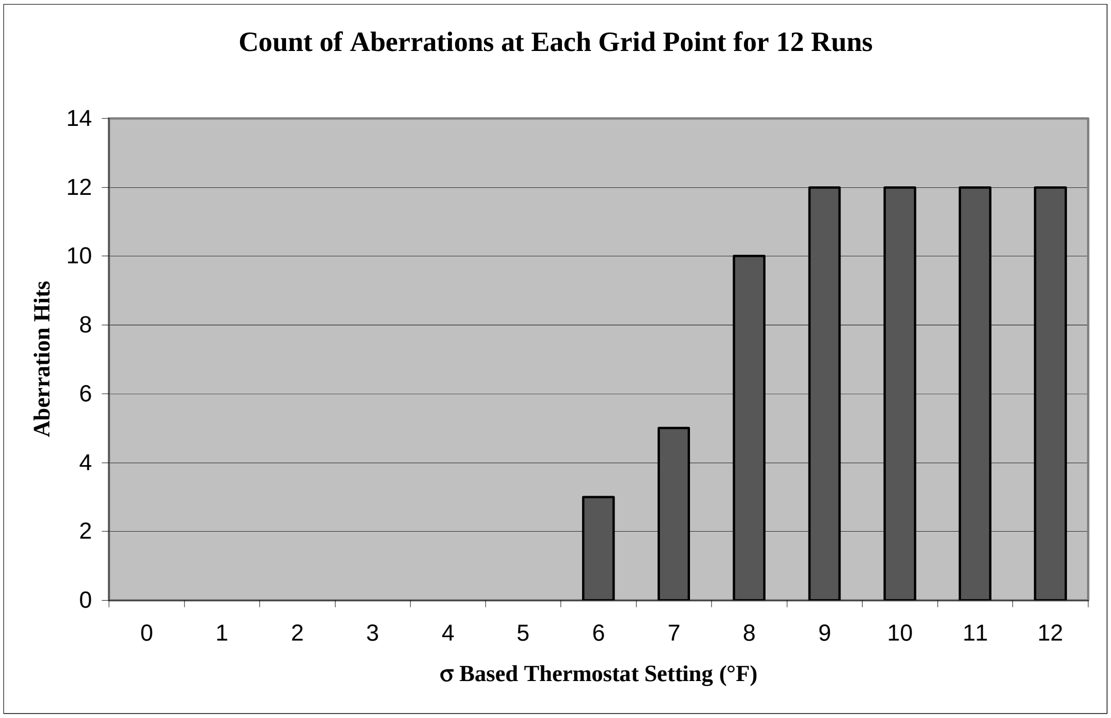

Figure 31 - Aberrations Present at each Grid Point (Current Practice)

The first instance of running into an anomaly occurs at 6 °F. Therefore, this value acts as the threshold for the range of magnitudes over which both thermostat settings are to be compared. The data shown in Fig. 31 is based on 24-hour complaint recovery periods. Running the same diagnostic test for the new strategy across the same set of magnitudes resulted in no aberrations at all for 10 runs, due to the fact the PI gains were tuned for a quick response to setpoint changes. Using sluggish PI gains will result in this aberration appearing at higher magnitudes. For the new strategy, the boundary for the range of durations setting out the grid will be 24 hours, due to it being the first instance of CRP values indicating aberrations. This seems apropos, because setback durations are on the same general order of magnitude as complaint recovery periods.

4.2 Optimization of New Strategy

Resultant performance metrics, shown both as a 3-D plot and an overhead contour view can be obtained by running the new thermostat setting policy over a grid with the ranges as described previously. The annual complaint metric shown in Fig. 32 can be read by the color codes for both views. The darker blue colors denote “valleys”, or areas where local or global minima exist. Lighter colors such as light blue, cyan, light green & brown are area of “medium” height, and dark red areas denote “peaks”, or areas where local or global maxima exist. The 3-D plots are combined surface/contour plots that illustrate the 2-D projection of the 3-D surface on the same plot, outlining the major features of the surface. The optimization objective is to find a global minimum, so the regions of interest are the dark blue areas.

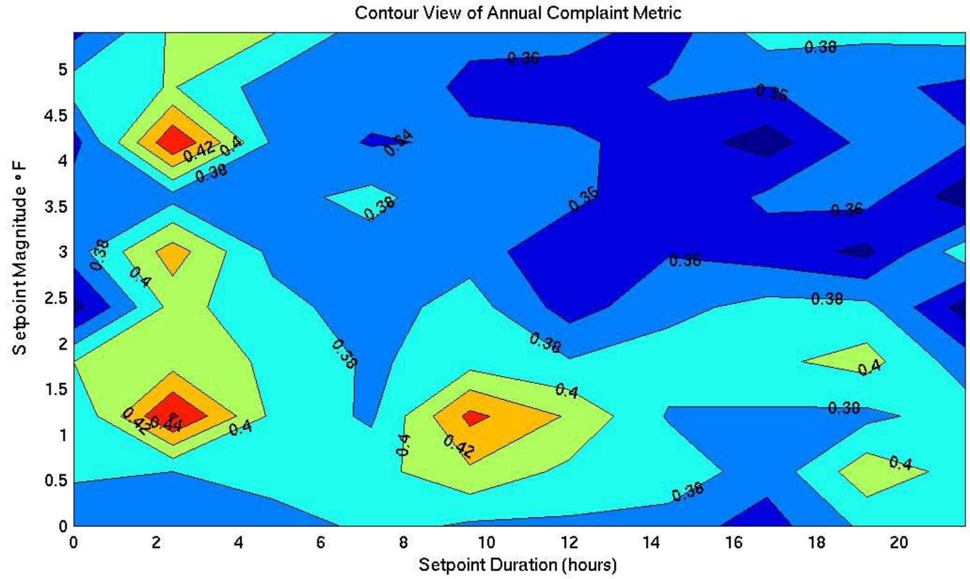

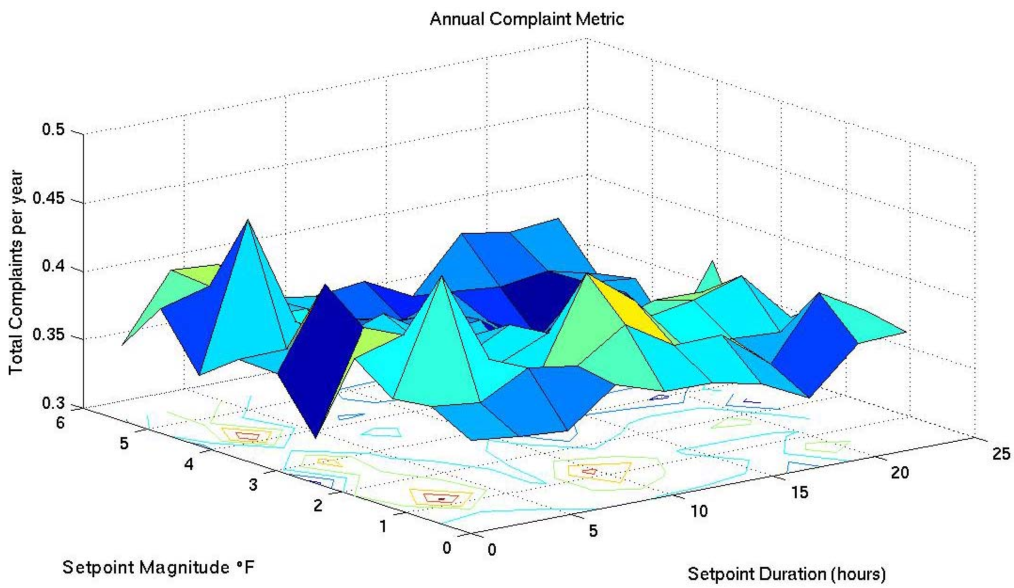

Figure 32 - Annual Complaint Metric

The second metric is the average complaint recovery period metric, as shown in Fig. 33:

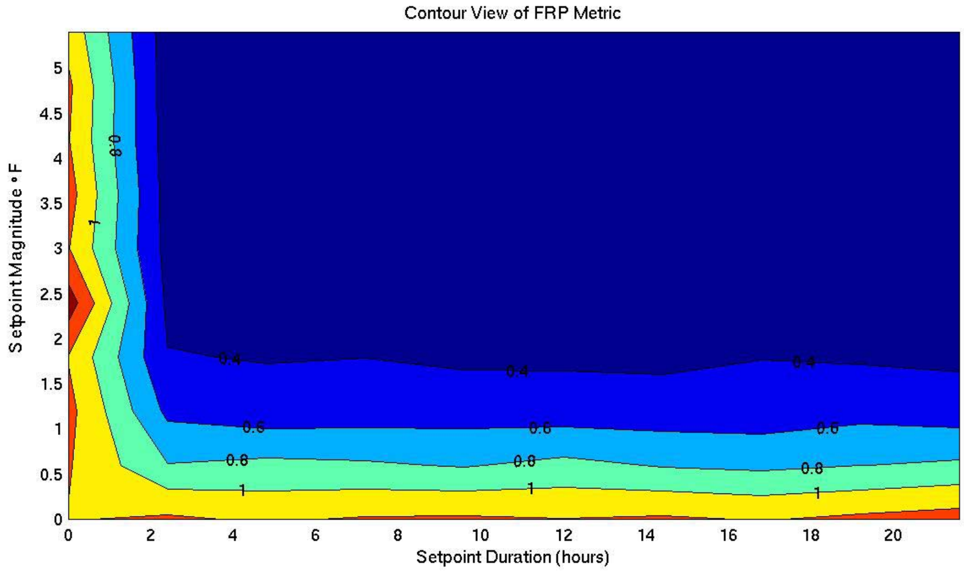

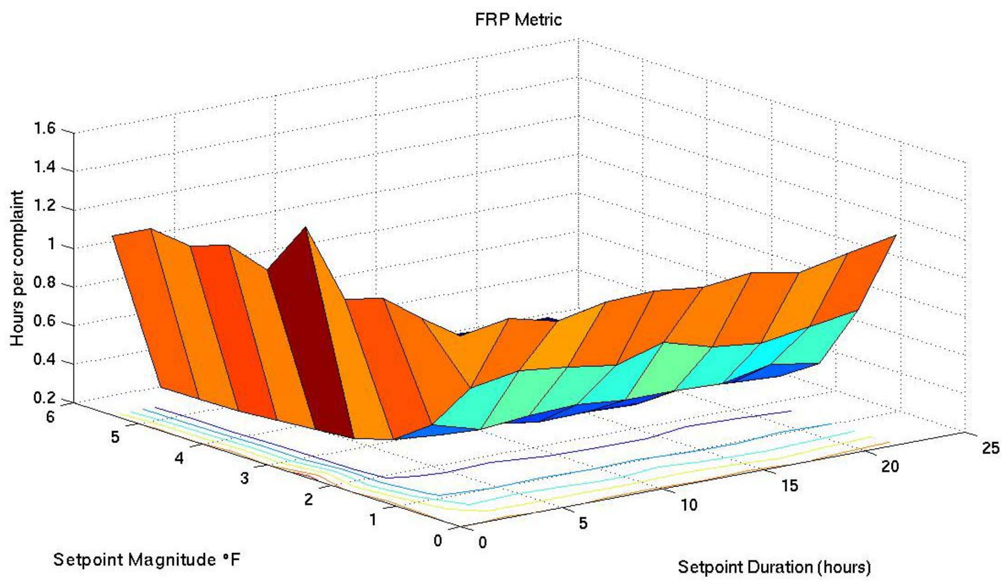

Figure 33 - Average Complaint Recovery Period Metric

As seen in Figs. 32 & 33, an initial hypothesis can be made as to what these results imply, prior to performing a formal optimization. The annual complaint metric grid in Fig. 32 is quite noisy, even though simulation to obtain these results was executed to a 5% maximum noise tolerance criteria, taking the computer equivalent runtime of about 1500 years. It appears that there may be a minimum value for the annual complaint metric at a high duration around 16 – 20 hours, and a magnitude of about 4 °F. However, the results are too noisy to say anything conclusive without extensive verification by running numerous ensembles. The most interesting observation is in the complaint recovery period result shown in Fig. 33. The FRP grid seems to set lower bounds on the magnitude & duration of about 2 °F and 2 hours, respectively. This indicates there must be a thermostat setting policy of some sort in response to complaints for a finite duration. Even though the final policy parameters are not well defined, it appears that responding to complaints with a finite thermostat change of at least 2 °F is required to minimize the performance metrics. It also is evident that parametrizing the policy by duration can be potentially advantageous when waiting at least 2 hours prior to setback.

4.2.1 Hybridization

A slightly more formal approach can be used to summarize the final results in developing the new thermostat setting policy. A simple method is to use a hybrid result of optimizing both metrics independently, using the following formula:

$$H = \rho R + (1-\rho) P \frac{R_{min}}{P_{min}}$$

-or-

$$H = \rho R \frac{P_{min}}{R_{min}} + (1-\rho) P$$

Equation 60 - Hybridization Formulae

In Eqn. 60, H = the final value of the hybridized performance metric, R = the average complaint recovery period metric, P = the annual complaint metric, and ρ = the percentage weight assigned to the average CRP metric. The remaining portion (1-ρ) is assigned to the annual complaint metric. Normalization is needed in order to obtain the final hybridized metric in terms of the units used by either metric. The optimized values of both the average CRP & the annual complaint metric, R_{min}, and P_{min}, respectively, are needed to perform this normalization. The basis for the final hybridized performance metric is left as a simulation parameter to be selected by the user, hence the two different formulas of Eqn. 60. Results for ρ = 0.5 (50 %) and ρ = 0.1 (10 %), using the CRP metric as the illustrative example are shown in Fig. 34:

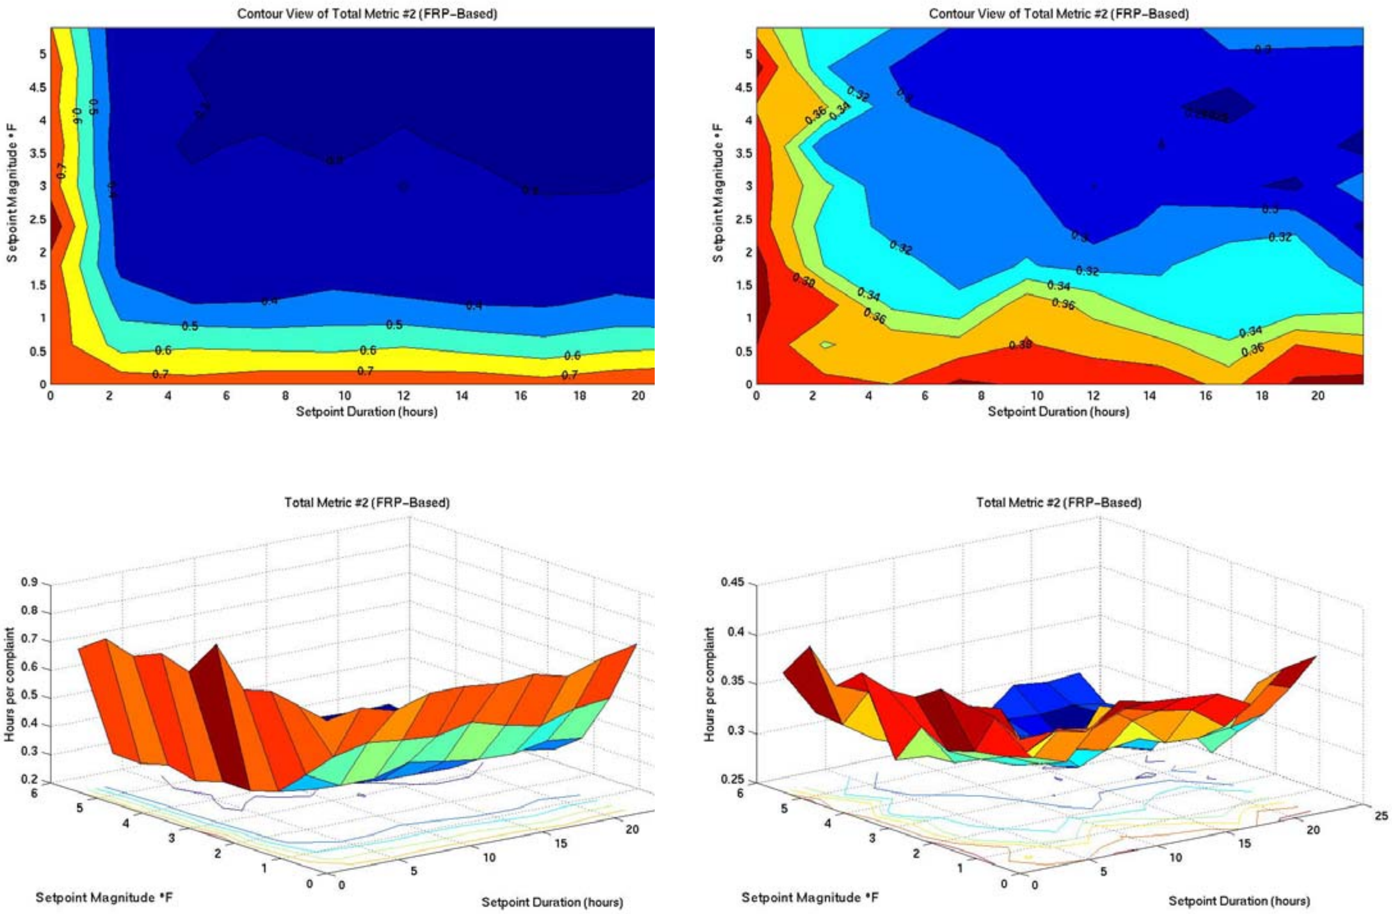

ρ = 0.5 ρ = 0.1

Figure 34 - Hybrid Comparison

Fig. 34 shows how using different values for ρ lead to different surface shapes used in optimization. When using higher values of ρ, the surface shape appears more like the FRP metric surface (Fig. 33), as opposed to using lower values of ρ, in which case the surface shape appears more like the annual complaint metric surface (Fig. 32). Using the hybrid surface is not necessary for finding the actual values of the minimum metrics themselves, because the pure un-hybridized surfaces can be used for that. A more tangible cost associated with how to choose the value of ρ is based upon the complaint recovery period, relating directly to how quickly the building occupant's complaint can be responded to, i.e. customer needs. On the other hand, the annual complaint metric relates directly to cost expenditures. Therefore, the tradeoff in appropriately choosing the weighting variable, ρ is left as a simulation parameter for the user to select in terms of what is more important, cost or customer needs.

4.2.2 Available Methods

There are several methods available to perform the formal optimization for verification or contradiction of the initial hypothesis made on brief overview of the final results. The first obvious method available is for the user to view the contour plots and make an estimation as to where the minimum areas are by using the color-coding schemes. Either the closest grid point to the values selected will act as the minimizing magnitude & duration, or they will be used as the basis of an interpolation scheme. In this case a least squares method is used to obtain the parameters describing the plane formed by the four points closest to the values selected. These parameters and the minimizing magnitude & duration provided by the user are used to find the value of the performance metric. Another obvious method is a pure minimization, where the appropriate MATLAB® command finds the minimum value on the grid & it's associated minimizing magnitude & duration.

As an alternative to these two basic methods, slightly more sophisticated methods can be used to find the minimum performance metrics & their associated minimizing magnitudes & durations. One method is to fit the performance metric surfaces with analytically described convex surfaces. There are two different types of convex surfaces used to fit the shape of the surfaces as closely as possible. The first surface type used for fitting is a simple paraboloid, while the second type is more complex, involving sinusoidal, exponential, linear, square, & constant terms, as well as linear & nonlinear combinations of all of the above. The least squares method is used to find the parameters describing both surface types, and therefore these surfaces can be plotted as well, over the same magnitudes & durations as the original grid. Both fitted surfaces can be optimized by using either pure minimization or the MATLAB® optimization call based upon the Nelder-Mead simplex direct search method. The cost function in question is the analytical function describing the fitted surface in compact matrix form, rather than all of the terms being written out explicitly. The initial starting vector used in the search is the minimizing magnitude & duration found by using the pure minimum of the surface found by using the appropriate MATLAB® command. Graphs depicting fitted surfaces of the annual complaint metric, using both types of surfaces are shown in Fig. 35

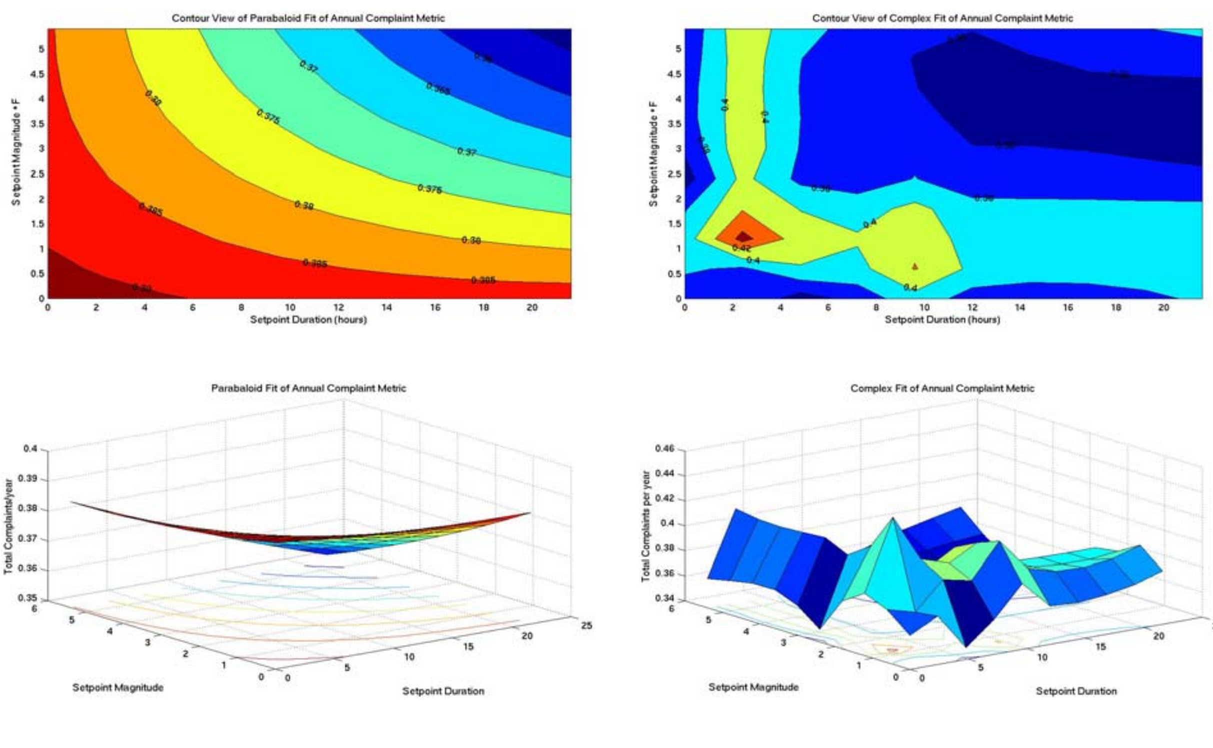

Paraboloid Fit Complex Fit

Figure 35 - Annual Complaint Metric Surface Fits

The general shape of the surface and the area appearing to contain the minimum annual complaint metric shown in Fig. 32 does not match the paraboloid fitted surface in Fig. 35. However, the complex fit appears to do a much better job of approximating the surface, and the area of the grid where the minimum occurs seems to coincide. Even so, using the ‘fminsearch’ method of finding the minimizing magnitude & duration is not capable of converging. Therefore the use of the surface fitting method is not always completely accurate. However, the pure minimum magnitude & duration of the complex fitted surface can be found easily, and is very close to the result obtained by using the previously described optimization methods. The quality of fit of the average CRP metric is shown in Fig. 36:

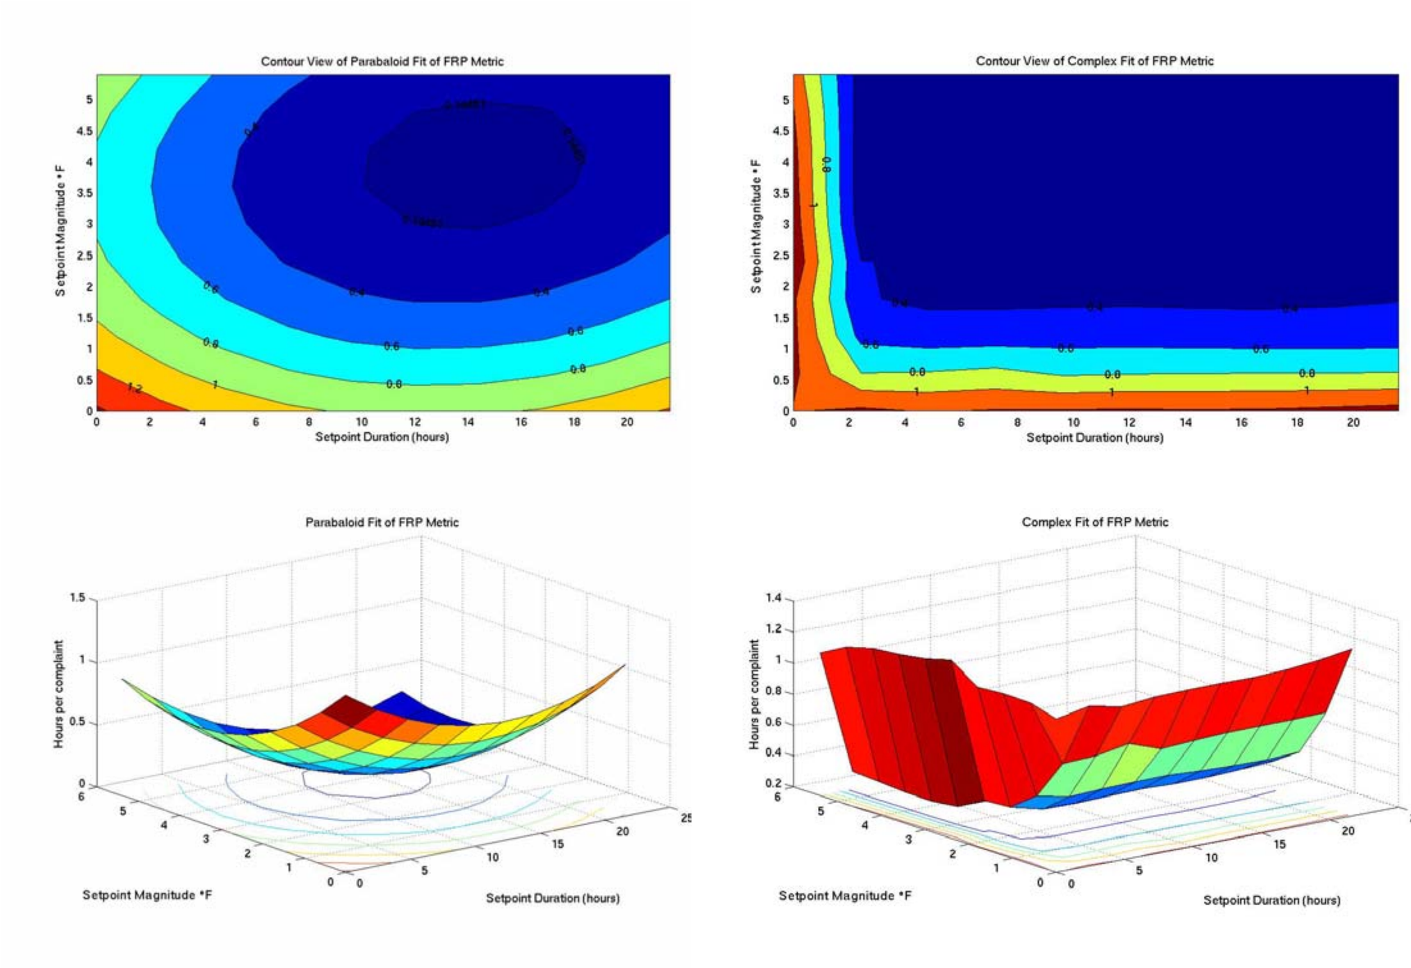

Paraboloid Fit Complex Fit

Figure 36 - Average CRP Metric Surface Fits

Even though applying the paraboloid surface fitting method to the average CRP metric results in a shape much different than the one in Fig. 33, it is obviously a convex surface which can easily be minimized. In fact, using the ‘fminsearch’ method to find the minimum performance metric & minimizing magnitude & duration converges to values close to those found when finding a pure minimum of the fitted surface. It also appears that the complex fit is quite good, and actually very similar to the shape of the surface seen in Fig. 33. Therefore, an analytical formula which accurately describes the shape of the simulated data can be obtained. Minimization of the complex fitted surface using either the ‘fminsearch’ or the pure minimization method will provide consistent results, giving the minimum performance metric & minimizing magnitude & duration.

Even though the complex and paraboloid fitted surfaces seem to be a feasible basis for optimization of the average CRP metric, there are alternative and superior strategies for optimization. The final method of optimization used for determining the new thermostat setpoint control policy is referred to as the “confidence interval” method. This method entails using the residual statistical uncertainty still present in the performance metrics after using the method described in Sec 3.4.2 to reduce it. As before, the statistical uncertainty of the annual complaint metric can be represented by the Poisson distribution. However, because the annual complaint metric is normalized, statistical uncertainty associated with each point in the grid must be obtained by multiplying the entire grid by the total simulation time length. This yields the number of complaints, which is required for use with the event-based Poisson distribution. Because the standard deviation of the Poisson distribution is equal to the square root of the number of complaints generated at each grid point, it’s easy to form a confidence interval associated with each grid point. This illustrates the confidence with which the metric can take on a given value in the interval given by: $[\lambda - \sqrt{\lambda}, \lambda + \sqrt{\lambda}]$, where λ = the number of complaints generated per grid point. There are some fundamental mathematical limitations on the values that can be associated with λ, due to the fact that square root of a value < 1 is always bigger than the value itself. Hence, values on the interval of $0 < \lambda < 1$ result in a negative lower bound on the confidence interval. However, because the complaint metric is based on integer-level complaints and not annualized complaints, it’s impossible for this to happen.

A region of minimality can be found clustered around the grid point associated with the pure minimum value of the complaint metric. It is defined by performing a radial search spiraling outwards, whose final terminating radius is left as an optimization parameter specified by the user. Therefore, concentric squares of increasing radius are formed around the initial minimum grid point, and each grid point is tested for confidence interval overlap. All grid points in the radial search with confidence intervals overlapping with the confidence interval of the initial minimum grid point are identified. These points form the region of minimality, and the centroid of this region acts as the new minimum. The idea behind this method is to find an area of minimality, rather than just relying on the pure minimum. If the stochastic uncertainty is large enough, there may be an isolated grid point that turns out to be the pure minimum, but is well below the value of grid points directly adjacent to it. Using a method where regions of minimality are dictated by several adjacent overlapping confidence intervals will tend to avoid isolated grid points being chosen as the minimum. An illustration of a side view of the grid displaying overlapping confidence intervals is shown in Fig. 37:

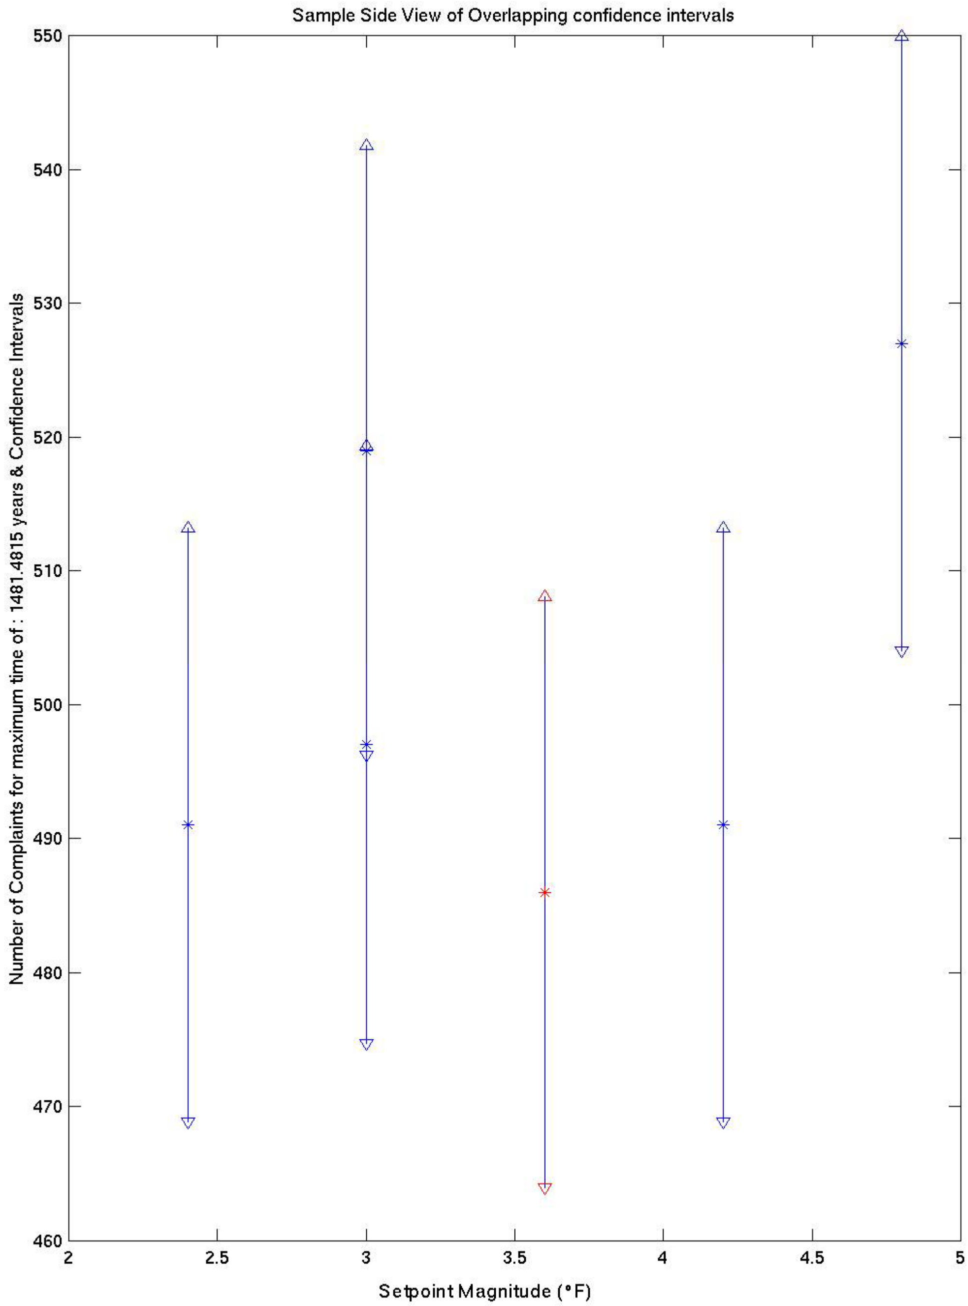

Figure 37 - Confidence Interval Overlap Side View

The confidence interval associated with the initial minimum grid point is denoted by the red arrows on either end of the interval, and the red asterisk denotes the center of the confidence interval, which is the actual value at that grid point. All overlapping adjacent confidence intervals are within a radius of 3 grid points away, which was the termination radius set for this particular run. From the side, it appears that at the row corresponding to 3 °F, there are actually two overlapping confidence intervals. The blue tips denote the boundaries of the each adjacent confidence interval, and the blue asterisks denote the center. The grid of overlapping confidence intervals from the top is as follows:

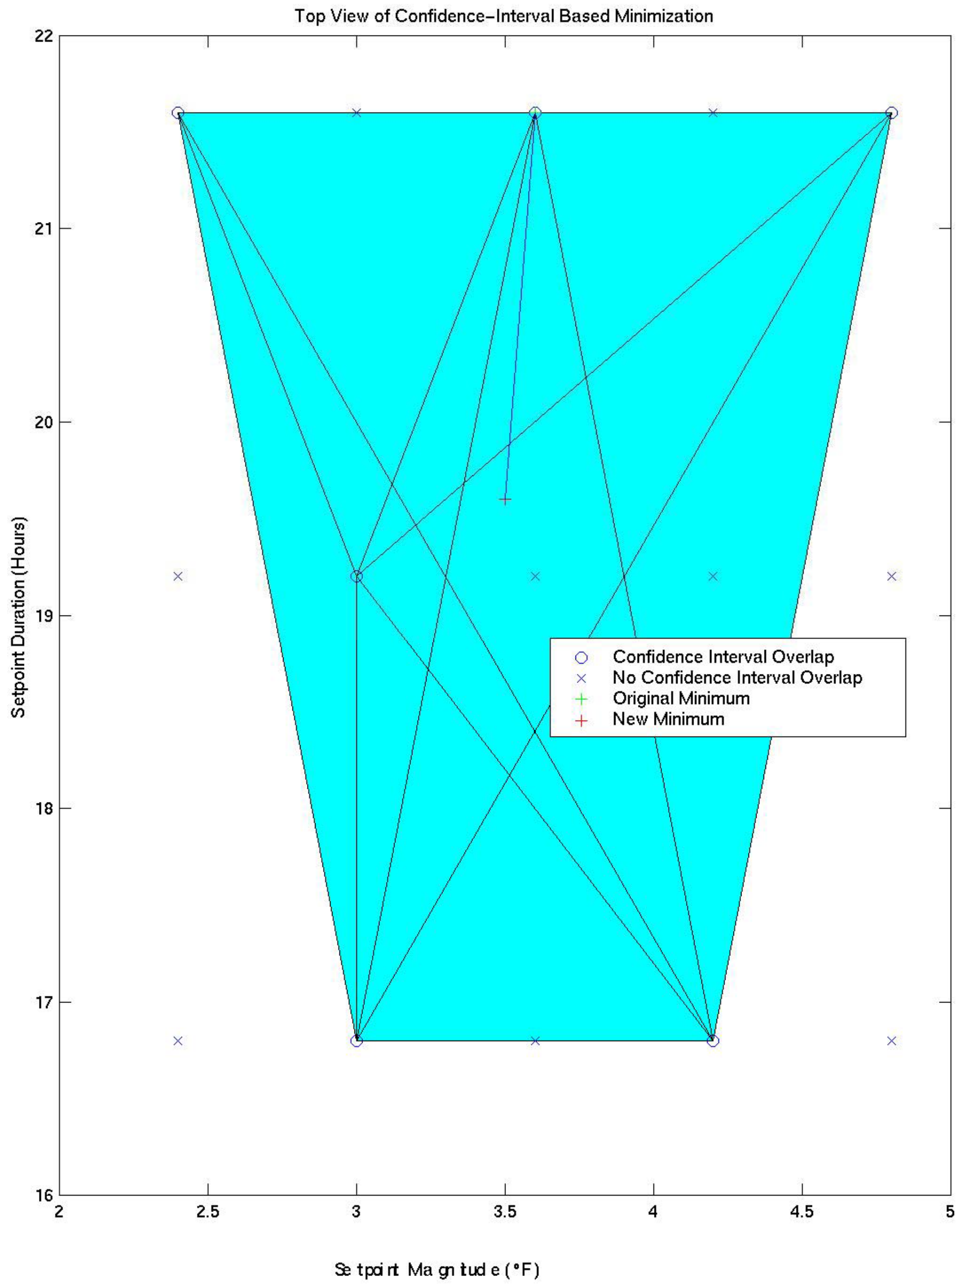

Figure 38 - Confidence Interval Overlap Top View

As shown on the legend in Fig. 38, the blue o denotes overlapping confidence intervals, and the blue x marks none. The green + represents the initial minimum grid point, and the red + represents the final centroid of the region of minimality formed by the grid points with overlapping confidence intervals. The blue line traces how the initial minimum moved to the new location at the centroid. The cyan colored centroid has several black lines that are an artifact of forming the patched centroid area in color. Because the new location of the minimum is not often associated with a particular grid point, the closest grid point or linear interpolation can be used to determine the minimum value.

A summary of the optimization methods used to obtain the final results, giving the new strategy's thermostat setting control policy, is provided in the following table:

	Minimum Complaints	**Minimum CRP**	**Thermostat Setting Policy**
Method	Confidence Interval	Pure Minimization	Confidence Interval
From	1481.5 yrs of complaints	1481.5 yrs of complaints	50/50 Hybrid Surface (Complaint-Based)
Complaint Value	0.3473 comp/yr	0.30359 hrs/comp	3.5 °F for 19.6 hrs.
CRP Value	0.38407 comp/yr	0.28042 hrs/comp	5.4 °F for 19.2 hrs.
Hybrid Value	0.33531 comp/yr	0.2887 hrs/comp	4.2 °F for 16.8 hrs.

Table 9 - Optimization Results Summary

The "Minimum Complaints" column refers to the optimization performed strictly on the complaint metric surface shown in Fig. 32, and the "Minimum CRP" column refers to the optimization performed on the CRP metric shown in Fig. 33. The last column refers to the optimization performed on a hybridized surface, similar to the ones shown in Fig. 34. The first row in the table shows the method used to optimize each of the three surfaces, and the second row shows the source of the data being optimized. The complaint & CRP surfaces each come from the simulation data, as shown above corresponding to 1481.5 years worth of complaints. The hybrid surface is based on a 50/50 balance between the two performance metrics, meaning that cost & customer needs are weighted equally in determining the thermostat setting control policy. The last three rows refer to the minimum values for the complaint, CRP & hybrid surfaces, using the minimizing magnitude & duration for the surface corresponding to the column. The actual minimizing magnitudes & durations themselves are shown as well. It appears that there is very little difference among the minimum values, regardless of which surface is being minimized over. Yet, discrepancies exist among the thermostat setting control policies as a result of the different methods used to determine them. Optimally the thermostat setting control policy should be dictated by the minimizing magnitude & duration values from the hybrid surface, and the pure un-hybridized surfaces for the actual values of the minimum metrics, which are shaded in the table above.

4.3 Comparison of Performance Metrics Plots - New Strategy vs. Industry Strategy

The current practice methods used in industry can now be compared to the new strategy developed from optimization, using the performance metrics as the basis for the comparison. As such, the strategy yielding the minimum performance metrics over a wide range of operating parameters can be determined. The graph shown in Fig. 39 demonstrates the results of the grid search for the annual complaint metric illustrating both the current practice and the new strategy. Although the current practice method is a statistically based thermostat setting policy, it can be compared on the same graph as the new strategy. The new strategy implements the thermostat setting formula $r = r \pm mag$, depending on whether responding to a hot or cold complaint, where mag is the deterministic magnitude of the setpoint change. The new strategy is parametrized by two variables, magnitude & duration, and the performance metrics are normally displayed on a 3-D plot. However, in order to compare the two setpoint strategies, only one row from the 3-D plot corresponding to the optimal duration of 16.8 hours across differing magnitudes is used. Note that the independent variable used for current practice is σ (statistical) in lieu of 'mag' (deterministic) for the new strategy, even though the x-axis in Fig. 39 does not make this distinction.

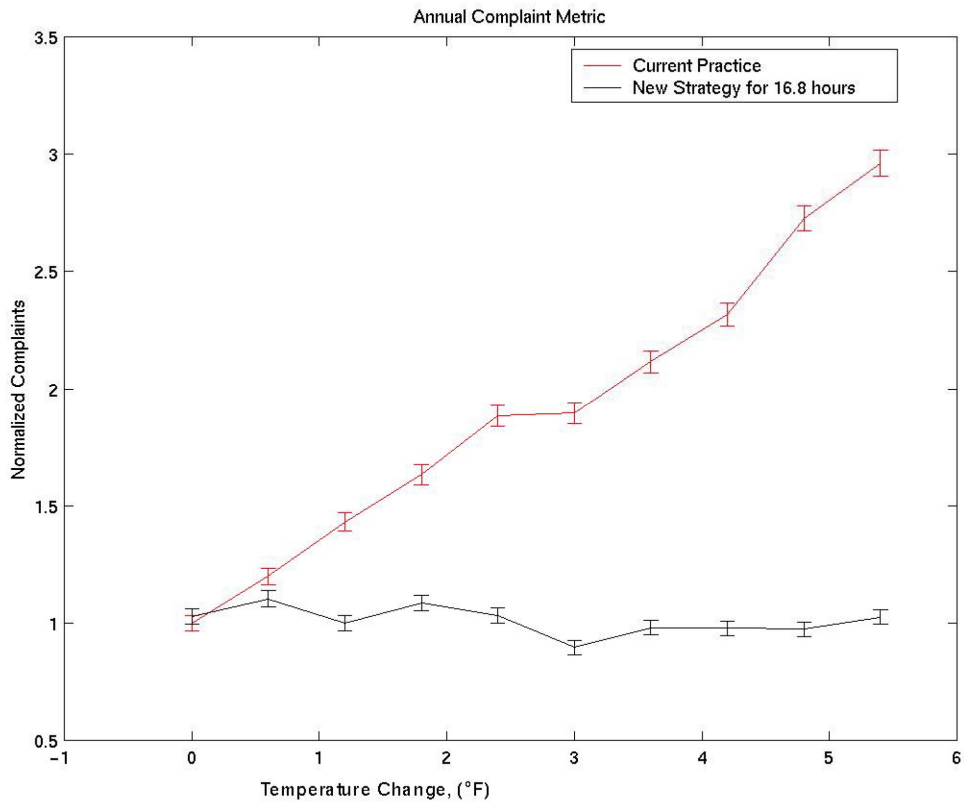

Figure 39 - Annual Complaint Metric Thermostat Setting Policy Comparison

It appears that even at magnitudes other than the 4.2 °F optimal magnitude, the new strategy consistently beats current practice methods in minimizing the annual complaint metric. Therefore the new strategy is superior over a reasonably wide operating range of magnitudes. Error bars are also shown in Fig. 39 to illustrate the statistical uncertainty still present in the performance metrics. Because the Poisson distribution has been used before to represent complaints as events, it can also be used as the limiting probability distribution representing the statistical uncertainty characterized by the error bars. From Eqn. 54, the cost function that results from using the Poisson parameter, λ, can be computed as follows: Cost = $\frac{\sigma[x]}{E[x]} = \frac{\sqrt{\lambda}}{\lambda} = \frac{1}{\sqrt{\lambda}}$. Because the expected value, E[x] is equivalent to the mean value, μ, of the random process, then the cost can be rewritten as follows: Cost = $\frac{\sigma}{\mu} = \frac{1}{\sqrt{\lambda}} \Rightarrow \sigma = \frac{\mu}{\sqrt{\lambda}}$. σ represents the standard deviation of the Poisson distribution. For good measure, the value of 4σ will be used as the width of the error bars shown in Fig. 39. The value of μ is equivalent to the running time average of the number of complaints. Using an approximation to Eqn. 53: $\mu \approx \lambda / L$, where L = total simulation runtime. Hence, the value of σ becomes: $\sigma \approx \sqrt{\lambda} / L$, and the formula for computing the width of the error bars becomes

$4\sigma \approx \frac{4\sqrt{\lambda}}{L}$. If the annual complaint metric is approximately μ, then the formula can alternatively be given by $4\sigma \approx 4\sqrt{\frac{\mu}{L}}$, where μ is the mean number of complaints per unit time.

Even though the complaint metric data is normalized, the formula used to compute the error bars is still valid because it is fundamentally independent of simulation time length. Simulation time length enters in the formula only from an approximation used to compute the mean. A second normalization is also performed by scaling all data shown on the plot in Fig. 39 by the first (0 °F change) grid point of current practice data. This ensures that the annual complaint metric at zero will always be unity for current practice, providing an easy basis for comparison of both thermostat setting control policies.

Similar observations can be made for the average complaint recovery period metric. The error bars shown in Fig. 40 for the average CRP metric are based upon the following equation:

$$\sigma = \sqrt{\frac{\sum_{i=1}^{N}(x_i - \overline{x})^2}{\lambda(N-1)}}$$

Equation 61 - Equation for Computing CRP Error Bars

In Eqn. 61, σ represents the standard deviation of the error per grid point, and N is the total number of recorded CRP's for each individual grid point. This data is collected independently of all of logged data such as the values of performance metric at each grid point, and the circular buffer of time series data & audit trail data logged for specific grid points selected by the user or at particular points of interest. Therefore, the value of x_i denotes a single CRP value, where $\overline{x}$ represents the mean value of all CRP values for that particular grid point. Without the extra λ in the denominator of Eqn. 63, it is simply the equation for the standard deviation. The square root of λ is added in the denominator of the formula to account for the CRP metrics being computed by normalizing the total simulation time spent in a complaint condition by the number of complaints for each grid point. Hence the metric is actually measured in hours per complaint, as opposed to hours. Therefore, the λ in the denominator accounts for this scaling when computing the error bars, whose total width is 4σ, where σ comes from Eqn. 61.

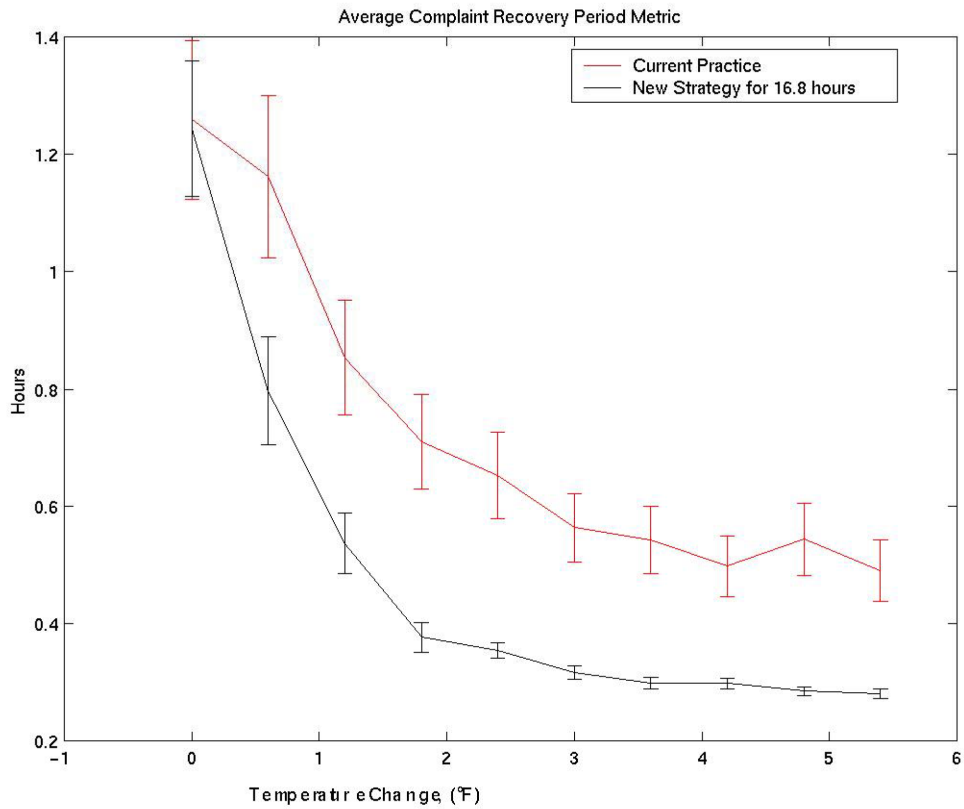

Figure 40 - Average CRP Metric Thermostat Setting Policy Comparison

Similar to the annual complaint metric, even at magnitudes other than the 4.2 °F optimal magnitude, the new strategy consistently beats current practice methods again, this time in minimizing the average CRP metric. Therefore the new strategy is still superior over the same wide operating range of magnitudes as before with the annual complaint metric. To summarize, it appears that resetting the thermostat to a nominal value after some finite duration is clearly a better thermostat setting control policy than a randomized, persistent setpoint change policy as is performed in current practice.

5 Discussion/Conclusions

5.1 Practical Usage and Interpretation of Results

From the results presented, it is evident that the new strategy will yield fewer complaints and shorten the complaint recovery period by using an optimized response. In fact, if setpoint changes over 2 °F are made & reset after over 2 hours, the dollar amount spent in response to complaints by using the new strategy can potentially be cut in half. Additionally, the shortened complaint recovery period can similarly be cut in half. Therefore, it is obviously cost effective as well as beneficial to improving building occupant satisfaction to make the second trip and readjust the thermostat setting to a nominal value. The large value of 16.8 hours for the setback duration portion of the optimized thermostat setting policy may be due to the fact that average CRP values are much smaller. Hence this long time may be optimal for "CRP avoidance", in order to ensure that the transient associated with the setpoint change in no way interferes with or is correlated to the complaint recovery period. This optimization procedure used to derive the thermostat setting policy also had no formal constraints other than those imposed by the range of the grid. Had there been specific constraints on the thermostat setting policy parameters, there might have been an empirical result with a more sound theoretical explanation. In fact, the suggestion of CRP avoidance is pure speculation because the noise still present in the hybrid surface is still too high to draw any definitive conclusions from.

The value of 4.2 °F found for the optimal thermostat setting magnitude also lacks a sound theoretical basis. It might be found that no local minima exist had the grid been run to values over 6 °F. However, a global minimum might exist at the magnitude edge of the redefined grid. In this case a large setpoint change, or "burst" response in magnitude would be most appropriate, providing the maximum temperature change for quick relief. The fact that the optimization value is 4.2 °F as opposed to the maximum possible magnitude on the border of the grid may have simply been an artifact of the statistical error still present in the results. The burst response interpretation is related to the fact that again, no formal constraints for the setpoint magnitude exist other than the range itself, hence a duration-only result emerges. All of these theories are still only speculations based upon optimization of the hybrid surface, which includes the statistical noise obviously present in the annual complaint metric. Therefore, a thermostat setting policy can be derived almost completely from the complaint recovery period metric, providing lower bounds on the thermostat setting parameters. The main usefulness of the hybrid surface optimization comes into play for projecting potential cost savings & improvements in response time in comparing the two thermostat setting policies. The bottom line is that sending facility operators back to readjust the thermostat setting to a nominal value is worthwhile, regardless of the complaint condition. This is the case assuming the thermostat setting policy is being implemented manually, otherwise a direct digital control system can be used to perform this action.

5.2 Future Work & Improvements

The work discussed shows that there are significant impacts to cost savings and improvement of customer satisfaction when using alternative yet optimal thermostat setpoint strategies developed from simple optimization methods. However, only the tip of the iceberg has been touched in terms in the research that needs to be performed. Integration of this type of an optimization into the current practice of building operations needs further investigation. As such, there are several points to make in terms of improving some of the methods used to arrive at the results presented in this report, as well as looking at some additional objectives.

5.2.1 Improved Modeling

Some of the modeling used as the basis of the research presented in this report were founded on simplified modeling conditions, which do not fully depict the conditions that might exist in reality. For example, some of the complex nonlinearities that exist in terms of the pressure dynamics of the room might need to be introduced, as well as looking at the interconnection of several rooms to form an actual building. Solar radiation through windows as well as other disturbances due to open doors, etc. might also be modeled. Furthermore, the actuator being used in the model is a simple air-handling unit that performs the function of heating and cooling. This actuator has no practical limits set on it as implemented currently; hence it is an infinite actuator. There should actually be saturation limits set on the heating & cooling capacity of this air-handling unit. A more accurate simulation of the hours of operation of the building is in order to delineate the actual times at which the most complaints might occur, related to the building HVAC system's startup & shutdown. The hot & cold complaint levels need to exhibit more realistic complaint behavior and associated delays. As such, the output of a coloring filer with relative degree of two may not be the best possible choice to simulate these levels. To determine realistic complaint behavior existing in practice, more extensive system identification & verification of the complaint model needs to be performed.

5.2.2 Additional Metrics, Optimization Schemes & Database Tie-In

The metrics currently being measured are just two possible metrics that might be of interest when keeping in mind how many different parameters in building operations exist in determining optimal performance. Some others might include complaint intensity, and energy used by the heating & cooling equipment. Complaint intensity might be measured in simulation as the area formed between the building & complaint temperatures when in a complaint condition. Equipment used for heating & cooling is just a single air-handling unit for this particular study. The supply & exhaust fans shown in Fig. 9 were not modeled. In reality, energy efficiency is a primary concern and becoming more important as management of today's energy resources heads towards deregulation. Therefore, all energy-consuming components of an HVAC system resulting in heating and cooling of a space, such as supply, return & exhaust air fans, compressors, air handling units, heat pumps, etc., should be included in this metric.

The optimization schemes used in this report vary tremendously. To tune the PI gains, a MATLAB® optimization call based upon the Nelder-Mead simplex (direct search) method was used. Prior to optimizing the performance metrics, an ad-hoc alternative to conventional methods was used to reduce the statistical uncertainty associated with them, which in itself could be cast as closed-form solution to an optimization problem. Finally, a number of different optimization methods, two trivial and two ad-hoc but more complicated methods were described & used to find the minimum performance metrics. It is possible that all of the cost functions requiring optimization could have used the MATLAB® optimization call based upon the Nelder-Mead simplex (direct search) method. For reduction of variance in the performance metrics, an actual Monte-Carlo (Vegas) method or even a purely analytical method may have been used. Furthermore, in order to determine the best thermostat setpoint strategy, a grid search constrained to meet certain parametric specifications relating to energy efficiency & direct implementation with a DDC-ready HVAC system is a distinct possibility for implementation. Perhaps it could even be constrained to search over only a framework of pragmatic thermostat setting policies to be implemented manually by facility operators, such as developing two separate thermostat setting policies: one for hot complaints, another for cold.

In addition to all the improvements listed thus far, there are a few final important ones to consider. They are all related to the fact that for this project, only thermal sensation complaints in a no-fault scenario are being considered. Hence, the HVAC system is operating properly & there are no latent faults that exist with any part of the system; equipment, ductwork, or otherwise. This is not always the case. In fact, as cited in the introduction, only 75% of all environmental complaints recorded in buildings are hot & cold complaints, as opposed to complaints about factors such as humidity, air circulation, etc. Furthermore, just 40% of the thermal sensation complaints happen when there are no faults in any of the HVAC systems servicing the spaces where building occupants reside. Therefore, future work may broaden the class of complaints being investigated to include not only no-fault thermal sensation complaints, but also thermal sensation complaints occurring under all conditions. Expert systems & fault detection logic could be used to examine these issues in a real-time, adaptive, online optimization procedure that determines the causal & probabilistic relationships between complaints & faults. Therefore, computation of optimal thermostat setpoints would proceed more intelligently with a lot more sensitivity to factors such as when preventive & corrective maintenance are performed on particular HVAC components. Using work request information associated with maintenance logs requires investigation of all connections relating this project to the bigger picture of building operations. This was illustrated in the introduction, tying in information from databases that serve maintenance management systems. Therefore, the complete realization of the interaction among building occupants, the HVAC systems & facility operators that serve them, and the facility managers that oversee such operations can be used to test these theories. They can be put into practice to counteract, predict, & respond efficiently to real complaints, so that real results, real cost savings & real improvements in customer satisfaction can be achieved.

6 Appendix

6.1 References

Auslander, David M. “Mechatronics : A Design and Implementation Methodology for Real Time Control Software.” U of California Berkeley, Department of Mechanical Engineering, 1997.

Dorf, Richard C., and Robert H. Bishop. Modern Control Systems. 8th ed. Menlo Park: Addison-Wesley, 1998.

Federspiel, Clifford C. “Predicting the Frequency of Hot and Cold Complaints in Buildings.” *1998 ACEEE Summer Study on Energy Efficiency in Buildings*, 1988.

--- “Statistical Analysis of Unsolicited Thermal Sensation Complaints in Commercial Buildings.” ASHRAE Transactions, *Proceedings of the 1998 ASHRAE Winter Meeting, Part (2) of 2, San Francisco, CA*, v104.1B (1998) : 912-923.

Fisher, David C., et al. Personal Interview. 12 February 1999.

Gelb, Arthur, et al., eds. Applied Optimal Estimation, Cambridge, Mass; London: MIT Press, 1974.

Haley, Roy. Personal Interview. 27 May 1999.

Incropera, Frank P., and David P. DeWitt. Fundamentals of Heat and Mass Transfer. 3rd ed. New York: John Wiley & Sons, 1990.

Newland, David E. An Introduction to Random Vibrations and Spectral Analysis. 2nd ed. London; New York: Longman, 1984.

Rohsenow, Warren, James P. Hartnett, and Ejup N. Ganic, eds. Handbook of Heat Transfer Applications. 2nd ed. New York : McGraw-Hill, 1985.

Strang, Gilbert. Linear Algebra and its Applications. 3rd ed. San Diego: Harcourt, 1988.

Tomizuka, Masayoshi. “Class Notes for ME232 Advanced Control Systems I.” U of California Berkeley, Department of Mechanical Engineering, 1997.

--- “Class Notes for ME233 Advanced Control Systems II.” Lectures given by Roberto Horowitz. U of California Berkeley, Department of Mechanical Engineering, Spring 1999.

Press, William H., et al. Numerical Recipes in C : The Art of Scientific Computing. 2nd ed. Cambridge [Cambridgeshire]; New York: Cambridge University Press, 1992.

Seem, John Ervin. “Modeling of Heat Transfer in Buildings.” Ph.D. Thesis U of Wisconsin-Madison, 1987.

Smothers, Fredric J. Personal Interview. 2 June 1999.

6.2 MATLAB® & C Code Execution Instructions

File Included on Floppy Diskette	Type	Purpose
pitunetest.m	MATLAB® script m-file	Tunes PI & JCI gains
premetrics.m	MATLAB® script m-file	Elicits all simulation parameters from the user required to set up the new strategy grid & simulate all temperatures; Run prior to simulation
premetricssd.m	MATLAB® script m-file	Elicits all simulation parameters from the user required to set up the industry strategy grid & simulate all temperatures; Run prior to simulation
premetricsopt.m	MATLAB® script m-file	Elicits all simulation parameters from the user required to set up the nominal thermostat setting grid & simulate all temperatures; Run prior to simulation
mopt.c	C-file	Runs the new strategy grid simulation
moptsd.c	C-file	Runs the industry strategy grid simulation
moptopt.c	C-file	Runs the nominal thermostat setting grid
cmetricsnew.m	MATLAB® script m-file	Performs the offline optimization of the new strategy thermostat setting control policy; Run after simulation
cmetricsnewsd.m	MATLAB® script m-file	Performs the offline comparison & analysis of the new & industry strategy results; Run after simulation
cmetricsnewopt.m	MATLAB® script m-file	Performs the offline presentation of the results & optimization of the nominal thermostat setting grid simulation; Run after simulation

Notes:

1) Only the most important files are listed above. Supporting files are also included on the attached floppy. All files are zipped in a file called MSfiles.zip. Unzip all files to a common directory.

2) The simulation sequence begins with running an m-file prefixed with ‘pre’, followed by running the executable version of the respective C-file, and then the final step is to run the respective m-file prefixed by ‘cmetrics’. Within this sequence of steps there should be enough interactive dialog prompting the user for simulation parameters such that the process is self-explanatory.

3) There are text files created by the MATLAB® script m-files run prior to simulation. These text files must be located in the same directory as the compiled executable versions of the respective C-files. Similarly, the text files generated by the C-files must be located in directory that is part of the MATLAB® path.

4) Prior to running the cmetricsnewsd.m MATLAB® script m-file, the MATLAB® workspace resulting from running cmetricsnew.m must be saved as ‘bigdurnl.mat’, in a common directory.

6.3 Stationary Engineers- Physical Plant Campus Services Interview

Identified expert(s) & expertise domain : Campus Services Physical Plant Stationary Engineers (Service Technicians) that respond to emergency & service calls to temperature & other complaints in buildings within a zone (Zone 4 - on & around 14-18 buildings central to Koshland Hall) located on campus.

The following is a transcription of an interview with a group (from 3 members to 5 members at various times during the interview) of service technicians :

RM - Rodney Martin (me)
ST - Service Technician # 1, 2, 3, 4 (hard to keep track)

RM : Ok, so, yea actually I wasn't having a hard time finding you guys down here, you said it was down here in kind of a bunker type area........

ST : (Interrupts) Uh-huh.......If they ever do a fly over we're in good shape

RM/ST : (Laughing)

RM : Alright...I guess the main thing that I'm doing here is trying to get a general idea of hot & cold temperature complaints ......I guess you service like a certain zone on campus

ST : (Interrupts) Yea....18 buildings

RM: 18 buildings, OK..........I was just kinda wondering.....if the actual frequency of hot & cold temperature complaints that you get in buildings are .....in general more prevalent than any other types of calls that you get ....?

ST : We get a lot of hot & cold.....we get quite a few.....that's our bread & butter., We get lots of different calls as well..........

RM: So you get the whole mixture .....

ST : Yea the whole mixture.......

RM: Whatever is out there, you guys go & respond to ....Ok..........but you say hot & cold temperature complaints are pretty much your bread & butter you say ?

ST: Mmhmm.....Very common call......Definitely not "THE" call, but ....

RM: If I asked you an approximate percentage, would you be able to say...........just roughly ?

ST: Hot & cold complaints.....percentage of our workday ? I'm gonna guess about 60 %...........(Next guy) I would guess 20 %........(Next guy) I was going to say 10 - 20 %

RM: Ok....and that's per day you say.......

ST: On an average if you want to take over the course of a month.......probably somewhere along those lines.............It has a lot do with if...it's smoking hot out there...........If it's a nice even-tempered day, people tend not to call

RM: Alright now .....I guess when you get a hot & cold temperature complaint in, what kind of stuff do you record, if any............... ?

ST: When we get the call, our dispatcher takes the callers name, phone number, building, room number, the complaint, and then there's a service request number for tracking purposes that's assigned to that.

RM: Do you know if there's a central ...database... that all of this information is entered into ?

ST: That's what............ she will enter that into the database.........

RM : (Interrupts) Ok....so the dispatcher will do that then..........

ST: Right the dispatcher will do that.......Now we will get the name, phone number, building, room and complaint, we don't get the service request number until it gets generated on the ....(speaking to other ST)........Do you have any ?..............Ok, it'll look something like this....(shows to me)............We probably have one and we'll Xerox one for you.............

RM: Great........

ST: And you can have that one to take with you ....

RM: Perfect that'll work great..........so do you generally need to generate an actual service request number, and .....get it through the system and all that ?............Do you ever have a really quick service call, or is that more for minor-type work ?

ST: A lot of it depends on the situation......If I'm in a building and this rooms hot and I'm here working on it and somebody comes up and says 'Hey the room 2 doors down (is hot as well) ....If I'm not going to be spending a lot of time on that one room, I'll just tack the other one on it. But if this one's going to be involved, we'll separate it............if it's a quick & dirty one we'll just bury it on another call......Because all of the service requests are entered back on our time cards.

RM : I see....

ST: And there's no point in having somebody type something up for a 5 minute line item on a time card.

RM: So ...you use that to charge your time to....

ST: Yes...

RM: Ok...Alright, and so, that does go from the dispatch log into a database....... Do you like have a certain............I guess just from experience.........Do you have certain things that you look for immediately like on a hot or cold temperature complaint, that you do as part of like a diagnosis ..see OK what's the problem here....like is something broken , or ....

ST: Well it starts when you walk in the building.....if the room's.......hot or cold...............when you walk in a building, your first impression is...How does the building feel........as a whole ?.......And as I walk to the thing, if I'm seeing a little variation in temperature, Ok, it's clueing me in that it's probably just one room........But if I'm walking down the hall, and it's sweltering, and there's (more than) one person who's complaining, I know that I need to go look at system-wide things as opposed to individual rooms.

RM: I see...

ST: And it's pretty much just experience that helps me (determine) where to look first. There's not a set procedure that ....the first thing you do is.......I go to the thermostat, I check its calibration, and then I go to the reheat......Well....(thinks).....essentially I'm doing that , and I'm doing those steps.....(thinks again)........I'm skipping those steps because I'm seeing other things that are influencing my decision that (the problem is) more likely over here.

RM: What types of things..........Are there any types of rules that you.....I mean not something like that's on paper but just something that you know, kind of by experience that you can say that this is definitely what I'd do as a rule.......you know kind of like a rule of thumb.

ST : Yeah.....I mean ...ideally, you go to the source, and you see what's causing the room to be to hot....

1) Is there too much equipment in the room ?
2) Are they on the sunny side ?
3) Is the fan blowing ?
4) Is the fan blowing hot air ?

.... Things I'm going to see right in the room. We'll check the thermostat, we'll check the reheat, whether it's a reheat valve, or double duct. The heat source specific to that room.......if all that's working, now I'm going to move back to the fans. And like I say if I walk down the hall, and it's hot, that's telling me that I could probably not hit (individual rooms) first, let me go hit the fans, and get the fans tuned in right,.....and then I can come back to the (room), and do those other steps, and change the order around (when I do them).

RM : Ok....What's your most common...remedy,......What's' your most common thing that's wrong on a call that you have with a temperature complaint ? Is it like a thermostat that's screwed up, or an actual system-wide type of problem ?

ST: It depends on the building.. Some buildings we've had failures of the air system, so it's getting oil into the thermostats.....so we're getting thermostat failures. In other buildings, we might be having problems with some of the older controls that serve fans.

RM : Ok, I see so control systems usually are..... ?

ST : (Interrupts) .....95% of the hot/cold calls are control system problems.

RM : Ah !.......Whew ! ...that's good, ...OK...(scribbles furiously on paper)

ST: As opposed to.....whereas the other ones are......

RM: (Interrupts)..Would you say that's because they're ..... the older ....pneumatic type control systems ?

ST : Well most of our buildings (have) pneumatic thermostats. Some of the newer ones and the animal facilities are electronic........direct digital control.

RM: Right , right....

ST: (Referring to newer DDC systems) And which still has a component that you still have the same valve in the field, you still have the thermostat.........you're still using air to drive it, you're still using that **same** valve. They're a lot better !! (enthusiastically)....if they're initially set up and tuned correctly......

RM: Yea....yep....absolutely ! I hear ya there.

ST: Then you have more to troubleshoot. (this was an outside comment from another service tech)

RM/ST: (All agree on that valid point made)......

ST: One of the nice things about pneumatics is at least.......When I started ....there was no DDC on this campus......there were a few computer controlled fans for start & stop (operation). So I learned pneumatic controls......and pneumatics are nice because....there's air there, or there's **not** air there, and it's a real easy check, you pull the air line out and see if its blowing or if its not. Direct digital controls......either voltage or no voltage......., which means that I have to get an instrument to read .....what I've got....... what it's doing. So it was nice because I was able to learn heating & ventilation with the old hand drawn things. And as DDC (came along)..I knew ventilation (already), and it was just a matter of what's controlling it.

RM: I see...OK, ...do have any ....you talked about the "path" that you perform on a basic type of service call. Do you have any kind of prioritization scheme in doing that ?, I mean do you know, for example, if you get a hot call.....

ST : (Interrupts) If I get a hot call in a building and I don't like the building, I don't go to them first ?....

RM/ST : (Laughs)

RM: But I mean just as far as ....if you get a complaint, the complaint is actually logged, and they say this is like the intensity of it....Like if it's 90 degrees in there or something, or do you prioritize it in terms of say, this is a certain type of building, like you were mentioning the animal facilities.

ST : Yea, ...Animals take number # 1 priority. It doesn't matter if I'm working on the Chancellor's heat, if you get an animal room problem, the Chancellor can leave the building, but the animals can't. ...Basically we look at what's going to have the greatest programmatic impact.

RM : I see..OK

ST: And if it's (emergent) we don't ignore the guy whose got one office where its 90 degrees. It may be the type of thing where we go over there and make a temporary fix, we shut the heating valve off because we don't have time to troubleshoot the control system. The heat's going to cool down, and then we can move on to the other emergencies, and then come back to him later.

RM : I see...

ST: And when we were there we would have looked at what type of control systems manufacturer's there were so that we would be able to bring a direct replacement instead of doing a complete modification....changing brands.......it's not difficult, it's just a little more time consuming. The one thing around here that's pressing....it's time.

RM: Alrighty, ...Do you know as far as the response time is concerned, for example from the time that you get a call to the time that you're dispatched, to the time that you actually get the thing resolved, if that's measured in any way, like by the central dispatch area ?

ST: No,...., and I'll give you a prime example...We were walking into the life sciences addition the other day and we were over there to check something else, and we were walking down the hall, and the building manager said "Oh you're here already ?".........And we're like.. "What are you talking about ?" ....(and he said)..... "Oh I just called this call in...." .... "Oh no, we just happened to be walking down the hall...........What's the call ?" ... So we took care of the building, figured out what we had to do, and about an hour and a half later, we were dispatched the call....So, the dispatcher will hold things depending on how urgent they are. "A" priority is life or property threatening.....that's' 24-hr response. This is how the guidelines lay out (whether it meets the criteria for urgency) or not. And it's not cross-supported by any stretch of the imagination. For example, a hot call would be considered a "B" tag. It's discomfort, it's not....if your office is hot, it's just annoying to you. If your crickets in the room are hot, that's affecting your experiments, it's affecting property, that would bump it up to an "A", but if it's just a too hot or cold call, because of the staffing level, it can't become a "B" tag and it just ends up in our box, and we say OK we'll try and get to it I guess we were planning on being over there we'll catch that too.

RM: Ok..Alright, so you .....just....try to... kill two birds with one stone if you can...

ST: Yea, we try to be as efficient as possible.

RM : So the dispatcher doesn't necessarily have like a ....

ST: (Interrupts) They don't even know where we are...unless they just sent us on an animal run call then they know we're there. And they can tell the client at that point...They're over working on an animal room, when they get done, they can move on to it (client's problem), but ......it's not going to be an immediate response.

RM : Ok...Just a couple more....How long would you say typically, I mean I know you aren't measured, but would say that the average service call takes ?

ST: Typical one....half an hour....45 min, if we're just goofing off. And that's' from the minute you walk in the front door.

RM: Ok

ST: You might be across campus or several buildings away, which, I view it as when I walk in the front door, we're there to do the job then, but the travel time is hard to gauge. If you gave me the call this morning, and I didn't' get there until tomorrow afternoon..........

RM: (Interrupts) So it (must also be) a function of when the dispatcher gives you the call

ST: It's when we become aware of it, and what else has come up.

RM: Ok..alright.....

ST: It's an art not a science...

RM : Yea...yea I hear you , I mean I remember when I used to work with the Navy Public Works Center down in San Diego, there was the same type of operation almost.....Ok.........well let me kind of try to tell you what I'm trying to do........

ST: Ok...

RM: Just so you know for your information.......and let you go ahead and tell me what you think actually, in terms of this idea. Basically what it is , I'm doing research for hot and cold temperature complaints, and these are .....unsolicited complaints, in other words they're things that you don't go (ask)......... 'Are you hot are you cold ?......'

ST: Oh you learn never to do that because once you bring it to their attention that there is something wrong, they say "Oh let me think about that ...", and sure enough, the next day you will get a call from them.

RM : Yep.......So I'm basically taking that, taking these complaints and trying to use that in a kind of feedback to a system that would take the type of knowledge that you guys have in terms of what types of things that you look at when you get these complaints, maybe just observations that you guys make, like if there's something, coming out of this wall, or, something is obviously broken, this is what you do to remedy the problem. Just setting up a database of those types of rules, and then using that to figure out what the remedy is, and possibly using that as feedback into one of these control systems.............And that's the whole basic idea, to get the knowledge where it wouldn’t be just one person but possibly multiple people’s knowledge. And that’s this whole idea of knowledge engineering, where I go around talking to any experts and trying to see if it is feasible to get that type of knowledge in a database and working on it from that level............So that’s pretty much what I’m trying to do, and I was just wondering what you might think about something like that. You know, would it help you, or aid you in any way, maybe trying to diagnose problems at all, or ....I know..in reality, that type of thing is ..you know in day to day operations, may be more of a hindrance, but I guess from maybe from some point of view.....

ST: No, but, I see where you’re going cause it’s like to get to those generations you’ve got to be here, and move (there). Things get more and more sophisticated. Like I say, first computer control systems are starting & stopping fans, now we’re resetting the .....(tape ran out)

Basically, he went on to validate my idea & thought that it was a good one.......

Additionally, he sort of summarized his whole viewpoint on my interview as follows :

In the process of actually going out to respond to a hot or cold complaint call, it’s analogous to looking at a maze, because there’s a certain procedure, a certain number of steps that needs to take place before getting to your final destination, which is to fix the problem. You don’t know all of the steps that you might need to take ahead of time (again analogous to a maze), so therefore, you might find that you went here and you didn’t need to perform that step. You don’t know that every step you take is necessary to be done to get to the final goal. So it’s a maze in the sense that you might have to go cross the same path more than once, and it’s certainly not the most efficient way, but you know that eventually you’re going to get the problem solved. Therefore, it’s jut by experience that you build a map of this maze in your head, so you know what the procedural steps are that are to be performed.

6.4 Supervisor - Physical Plant Campus Services Interview

Interview for knowledge acquisition with Roy Haley, Zone 4 Maintenance Supervisor

Legend : ***RM:*** Rodney Martin, ***RH:*** Roy Haley

RM: Explains purpose of interview (hot & cold complaints), as well as master’s project to initiate dialogue

RH: A lot of our complaints (that) come in are, I want to say physically driven. Women have a higher complaint ratio than men, by far. Now I don’t know if that’s because the campus is, on the administrative side female-based or not, I don’t know what the ratio is. But many things seem to affect the women, and the surrounding temperatures. Menstrual period, I think is very possibly something you could put (as a factor) here. Weight, age, change of life, and because a lot of the spaces which were at one time designed to be open large spaces, and through the new modern quarters & such, you can equalize....we don’t have adequate airflows. Right now the problem you saw us working on when you came in is where they just did a very high dollar ...retrofit, and they forgot to add the heating hot water back in that they took out. But as far as complaints go, in my observances in (my) almost 10 years are, women have a much higher ratio than men, and a lot of times its because of the locations of the diffusers. A lot of situations we’ll find there will be two desks facing each together. A lady in University Hall, this was a problem, the diffuser was over the center of two desks, the lady on one side was a Caucasian lady, very heavy set. On the opposite side, was an Asian lady; probably didn’t weigh 90 pounds. One was burning all the time, and one was freezing all the time. And basically we ended up solving this problem by turning the whole register off, because there was no way I was going to make a divided register, turn half off and everything else. And the building manager just got so upset with these two women, bickering and fighting over this, that she just had me come in after the ladies weren’t there, and I just turned the register off, and to my knowledge it’s still turned off to this day.

RM: No complaints after that ?

RH : No complaints. And that's something else, if you have more than one worker, two or three, sharing a common thermostat, you're going to run into a lot of difficulties, so that's something you might want to think about. In a single cubicle, like I'm in here, in an enclosed room, and I'm the primary resident, I can do what I want with the temperature. If I'm using your type of 'mind control' program, that's fine, but if you get in an area where you have more than one occupant, who is going to be the prime leader of the control for the temperature ?

RM: Right.

RH: So that's something to think about

RM: OK

RH: Another big problem is...customers only hear what they want to hear. This is standard whether you're buying an automobile or a dishwasher, or heating/cooling complaints. We have the same ladies, and a couple men, every year during the summer who call in and complaint about being hot. They've probably been told 200 times, their building doesn't have air conditioning. Whatever the outside ambient air temperature is, is what the room temperature's going to be. And these are reoccurring, I don't whether it's a placebo, if it makes them feel better. They've got somebody they can call up and complain to. Another thing is we live in a coastal area, our mornings are notoriously chilly, and as the day warms up, and people don't dress for the morning hours. They'll come to work (with a) shirt maybe, and maybe a light tee shirt underneath it, and..... It's 50 degrees outside, and they get chilled in their transiting. Anyway, when they get here, the campus, most of it isn't set up like your home, where you have a register, and you've got this big blast of warm air flowing out right at you. We use a lot of indirect type heating, and we also follow the state law. And the state law, right now, says unless there are children, animals, or special conditions, no thermostat is supposed to be set above 68 degrees Fahrenheit. And we do fudge on that a lot, just to keep our clients happy.

RM: Can I interrupt you for a second ?, because I think we're at the point now where I can ask you some specific questions, that 68 degrees Fahrenheit would be a good starting point. Ok, so, basically, in response to getting a hot complaint, for example, I'm assuming there's a building coordinator who takes the complaints in and will report them to someone who actually responds to them.

RH: Right

RM: Do you actually change the setpoint of that system, in response to that complaint, and then if so, then how often do you do it ?

RH : OK, why don't we start.....Do you want to start with how we get to a complaint to my stationary engineers, do you know that process already ?

RM : I know that process

RH: Ok great...This is more or less a judgement call of the engineers. They have a more hands on or face to face involvement with the clients than I do. I normally deal with the building managers, and the associate chairs and the chairs, and people like that. I do talk to clients, and I do go smooth ruffled feathers, whatever, but I don't go on hot and cold calls, my stationary engineers do. It's kind of a judgement thing, one you (start to) get (a) repeat call(s), you kinda learn the lady......or the man. I don't want to be sexist or whatever. You learn the client, and sometimes you can just walk in and you can just fiddle with it, and say "Now are you comfortable ?",......... "Oh, yes so much more"...., and you may have done nothing. And it's just that you're appeasing them.

RM: Like a placebo

RH: Yes, and at other times we do fudge. Sometimes, a lady will say I've been really sick the last six weeks or whatever, can you warm my office up a little bit ? And the guys, most of them, will maybe bump it up by 3 or 2, which helps them out, plus it's also a mindset for them.

RM: Ok, so when I ask when the setpoint is changed, how much is it changed on average, and for how long, then really you're saying it depends upon the situation.

RH: It's a big variable.

RM: So, what strategy is used for changing the setpoint in response to complaints, is it more of an adhoc as opposed to actually a fixed kind of procedure that you have for changing the setpoint ?

RH: Yea, it's client based, a lot of it, and surrounding based. Client based sometimes is just...we were running real shorthanded, and to eliminate a lot of the nuisance calls, you raise the (thermo)stat a couple of degrees or whatever. And sometimes there's conditions, they may be doing some work down the hallway where they have a large window open, and they have more outside air in the building, so the specific room, maybe her door is open & closed a lot, there's more foot traffic.......

RM: Let me pose this situation: Would there or do you think there would ever be a case in which a person would complain if they're too hot, so that you change the setpoint down, for example from 72 to 68 degrees, and then after a certain period of time expires, you set it back to what the original setpoint was before you changed it (72) ?

RH: We don't have the manpower to do that. I mean, that would probably be a pretty good way to do it. If you knew you had a client that during the call, maybe was a little colder than the rest of their peers in the immediate area, and then maybe she's just more temperature sensitive, and then when summer got here, she's also more sensitive to the heat. But we don't have the people to go around and do that. So we're not so much proactive as reactive, from the management sense.

RM: So, I guess when you do, like if the setpoint is changed, from a certain value, I'm wondering if that change is actually logged anywhere, like in a database that you have where you store a bunch of complaints, anywhere at all ?

RH: We do have rooms that are very temperature sensitive, but these aren't occupied normally by clients. They're.....research, holding areas, lockboxes, etc., ..We do have some laboratories that are occupied that have temperature requirements. But those are pretty much fixed, and they're chart recorded, and there is a very defined window that we operate in. The client will tell us , in animal research this is very predominant, they'll give us a range, and sometimes its like less than one degree centigrade plus or minus, that they want these particular animals kept at, and we have to maintain that.

RM: It's a really tight tolerance on that.

RH: Yeah, some cases were extremely tight, but for human occupancy, we use your generic Johnson & Honeywell, Powerstats & controls.

RM: Direct digital controls. Yea, exactly, that's kind of the target of part of this (research) as well. Ok, so complaints in general, they're logged somewhere in the database.

RH: Right

RM: But do you distinguish the hot & cold complaints from the other complaints on the (log) ?

RH: Out database is set up by buildings and by the trades. So I can go in, and you can give me a building, and you can say, OK hot or cold, I'd look up under stationary engineer for that particular building. And I can also search by timeframe, so I have a few search engines, but our database system is very very old, in fact it's in the process of being upgraded now. Maybe the new one will give us that capability, but for now I have to attack it from a couple different angles, rather than just hot or cold.

RM: I remember he gave me, the last time I was here, I think I was talking to one of the stationary engineers, he gave me a sample....call.

RH: Service request

RM: Service request here, and it seems like there's like a certain field here in which it says description and action taken. OK like for example this one, it says Library is cold, so you know that was a cold complaint. Would that be the field where you would search for like the keyword, so to speak ?

RH: No, our primary search capability is this : Service request number. Our secondary search would be the building name, this four-letter code here represents. OK, another search field that we utilize would be stationary engineer, that's what that abbreviation is , STAT. Our program is very old, most of it was set up on a four digit.......

RM: (Interrupting). Oh ok, so those are your only three search fields.

RH: Well, actually, we can search by funding account, we can search by date, and that's about it I guess.

RM: Not really descriptions or anything like that ?

RH: No,.....

RM: Ok, there are ways to import this type of database system where you probably would have access to these other fields, but that's all, like you said eventually going to happen. All right, so I guess the next thing would be.........When you're logging these, these are just like the initial requests for this.

RH: Right

RM: If a service technician goes out and responds to one of these hot complaints, and he finds that as a result of whatever investigation he does, maybe beyond just the person being cold and changing the setpoint, finds that there's something wrong with the actual system. That problem with the system, is that logged in the database as well, or is there some kind of a record ?

RH: The procedure is the stationary engineer goes first on a hot/cold call, and he goes in and he is also the maintenance person that would perform the repair, replacement of the thermostat. Ok, so our engineer has a laser reading temperature sensor, so he doesn't have to wait for an analog type thermometer. He can come in and immediately zap it, and a note to point (out) is that a lot of people, especially the women, even have bought their own thermometers, and brought them in, and most of time they're way out of calibration. But you can't make them believe that........ "I just paid $3 for this thing, I know it works"........and the thing will be 10 degrees out of calibration. Ok, if it's a problem that he can readily identify, and it's in his area of repairs, if he has the parts available he can repair it right then, if he doesn't, ...he informs the client, ...comes back to the shop and gets the parts, whatever the case may be. Say it's a control valve that the thermostat is driving, that falls under the realm of steamfitters. So what he'll do, is he'll come back, he'll look at the service request. He'll come back, and he fills out a paper, which he gives to Chris. And Chris will type up another service request. It'll have the same SR#, but see where it says task number ?, It would be '02'. Now say that the steamfitter goes out there, and he looks at it, and he says 'wait a second, it's not the control valve, I don't think there's enough air coming from the whole unit'. So he may come back , and talk to Paul on the radio and say....And task '03' may go to an electrician, to check the motor rpm. Ok, so in other words you can build on these task numbers, but this SR# is assigned solely to Warren Hall, Room 42, the library is cold. (i.e. the initial request). Ok, so this number will remain constant, but the different numbers here, the task numbers, can be added to, and the new tasks like say it was '02' assigned to the steamfitter, this down here would say steamfitter then. And it's an assist to stationary engineer, check control valve, room cold.

RM: Great, so there sounds like then there is some type of search capability so that you can search by these two numbers.

RH: Right

RM: So that you can define kind of like a historical information (log) on the initiation of that original service request.

RH: Right

RM: Ok, excellent. So, I'm wondering, and I'm going to have to define some things for you.......I'm going to define a hard fault as opposed to a soft fault. And my definition of a hard fault is something that will affect the integrity of the system in such a way that if you change the setpoint, it doesn't matter, because something's broken in there, something's stuck.

RH: Oh, ok

RM: **.....**That's keeping it, so if you change it to 110 degrees, it's still going to be operating at whatever .....

RH: No control...

RM: No control.....And a soft fault, is more of something where it's like a miscalibration, like something's kind of over time, creeps to a certain level, just off ......that type of thing, So it doesn't really affect the integrity of the system.

RH: Right

RM: So, I guess based on that, just on your kind of subjective assessment, I guess, how often would you say that the reason for hot complaints are due to a hard fault.

RH: In this environment, in the collegiate environment, (considering) the damage done to the thermostats, I'll say it probably runs 60% soft, 40% hard.

RM: Ok. What about for a cold complaint ?

RH: Probably pretty much the same. Unfortunately, thermostats in a lot of the new electronic controls are made to work daily, up & down, up & down, up & down, by different operators, and here, we found that if we put institutional covers over the (thermo)stats, they rip those off so they can get to the (thermo)stats.

RM: Laughs. Yep...

RH: And...these (thermo)stats, they're fairly, they're not real expensive, they're $75 a piece, but they don't take abuse, and even though they have little set screws to lock the caps on, they pry the caps off, and then they bend the little (???) bars. We have..a lot of the control is found with the thermostat, you know the problems. Probably 60 %, and a lot of that is just through misabuse.

RM: I guess the next logical question would be how often are the reasons for these hot complaints or cold complaints due to no fault at all. In other words the guy goes out, and he changes the setpoint due to the response to that hot complaint, but there's really nothing wrong with the system.

RH: You mean the placebo type thing ?

RM: Exactly

RH: So you're just trying to satisfy the customer.

RM: What percentage of the time does that happen ?

RH: They probably do that maybe about 20% of the time.

RM: So I guess I'll have to normalize this a little bit, I'm trying to add all these three together.

RH: Oh, OK, If you want it to all add up, I think the 40% hard, and make the other two equal, 60 % total, the stronger on the soft, and then, why don't we say 40% 40% 20%, how's that ?

RM: Ok so, 40.....

RH: 40% hard, 40% soft, and 20% the placebo type thing

RM: Ok, excellent, and that's for both hot & cold.

RH: Yea

RM: OK, perfect. All right, so a couple more questions along those lines, which are, these are going to be a little bit more, kind of strange situations. How often would there be a hard fault, when there's no complaint at all. In other words, like, you don't get a complaint, but the guy runs into a problem just because maybe he's checking something else, and finds there's something wrong with the system, and is surprised that no ones complained, but no one has. That type of situation

RH : A lot of leaks are found that way. The system will still operate because your heating & cooling systems all are under constant make-up capability. And you might develop a small leak, and it won't affect the operation of the system where there's chilled water or heating water.................it doesn't affect system operation, but then you have to bring the system down actually to do the repair. That happens....normally ventilation problems can be identified either by temperature or flow (diagnostically). A big thing nowadays is people are putting heat sources, heat loading in the rooms that weren't designed to carry that much heat load. Computers, copiers, fax machines, you know they generate a lot of heat, and also they're breaking up the natural airflow in the rooms, by partitioning in cubicles. We don't have the proper (???).....you know the circulation of the air. But state law does mandate certain requirements, like all laboratories have to have 100% outside air, animal facilities are 100% outside air. Buildings have to have a minimum of 10% outside air at all times. I'm sure you're aware of all these.....There are occasions where....

RM: Is that rare, or what would be the percentage of that ?

RH: Oh, do we have to figure this into the 100%, or do we start over again ?

RM: No, this is another 100% again.

RH: Ok, I would say this is rather small actually, maybe 10-15 %, 10%.

RM: Ok that's for a hot complaint...Would you say about the same for a cold complaint then ?

RH: Yea. You've got to remember too a lot of state buildings aren't air-conditioned. A lot of the buildings here on the campus aren't. You know like the new Soda Hall is, but that's because it's a computer science building. A lot of the actual teaching-type buildings are not air-conditioned, so you'll normally have more hot calls, and they're seasonal, because people don't realize they're (the buildings) not air-conditioned, and they think that the air-conditioning has failed, when in reality they're just getting ambient air temperature. And in the winter naturally, you're going to get more cold calls, because people are chilled from outside....because (they think that) the calibration or whatever has failed.

RM: Ok,....so here's this one that's really weird...I guess it's going to have to be the other 80% but, I guess what's the probability that you don't have anything wrong with the system, given that no one is complaining. So, ...

RH: We have a preventive maintenance program in place, in all honesty, it's way undermanned, I mean we're undermanned, it's way behind, it's not being done properly, we all recognize it. The preventive maintenance program,(for example),.......calls out for the (stationary) engineer to like check the belts on the fans. We don't go as far as to check the control valve operation or anything like that. The filters are cycled through periodically, we get the size filters and their locations. We have a computer program that one of the guys made up, and we track them so we know when to order them, you know when to replace them. That's it, I mean we just don't have the manpower to walk around and look for things broken because we have things broken already........

RM: That's just PM (Planned/Preventive Maintenance) as opposed to CM (Corrective Maintenance).....

RH: Yea.

RM: So I guess those types of situation aren't really caught that often.

RH: No, to be honest with you, no.

RM: OK

RH: There's usually some kind of tattletale; the noise, loss of ventilation and a leak, something that (came) from the customer complaint will drive us to investigate the problem.

RM: Would you say that if no one complained, the probability that there is nothing wrong with the system is fairly high ?

RH: Yea, I mean because normally, the pumps, you know if you have a seal failure or a motor failure, or something like that, you're going to know about it. Visually, when the guy is walking through the machinery room, if you have a seal failure on the hot water circulating pump, yknow, see the water on the floor or whatever, or if you have a motor seize, or a coupling break; then you're not going to have any movement of water, and you're going to get a lot of cold calls. And also, almost all of our systems are pneumatic, or electro-pneumatic. (If) we lose an air compressor, the whole system goes down, and people will let you know about it in a hurry.

RM: So say about like 80 % then ?

RH: (Nods)

RM: That'll work. Ok just a couple more here, and then I should be finished, let me just see if this thing is still running......So, if there isn't any fault in the system, at all, ...(skips question because we covered it). Do you ever get a hot or cold complaint for which you don't do anything at all ? I mean not anything other than change the setpoint.

RH: Uh..yea, just to inform the client if they don't have air conditioning. I mean just from walking in, or sometimes you can walk in and the lady will be (submitting) a cold call, and she'll say, oh I'm freezing, and it's a shared space. It's not uncommon for other women to say "Well we're all burning up". OK, so that's when I go "Ladies, I'll tell you what you do, ..you have to work it out, because according to state law it's supposed to be 68 degrees, 68-70, we'll try to keep you comfortable, because I can't interfere here, yknow, and sometimes it takes their supervisor to get involved to quell things". There are times where we'll show up, actually make the customer happy, and leave, and never touch anything.

RM: Ok, great......I guess before I ask you the last question, I'm kind of just curious about the process. I've forgotten a little bit about it, but, there is like a building coordinator, like you said like a supervisor who would kind of take the complaints of the employees in the space ?

RH: It's supposed to work that way

RM: But it doesn't always ..

RH: More often not than does, unfortunately. Some buildings are very good about it, Valley Life Science building, LSA, are very very good about it. Some of the other buildings..not so good. People just pick up the phone, and y'know they've been around for a year or two or they ask the lady down the hall who's been here for ten years, and they give her Chris's phone number, who's our dispatcher. And they call in a room hot call, well, it's not our place to say you should go through the building manager. We'd like to be able to say that, but that's not good customer service. And then what happens if the building manager is gone for two weeks, or what happens if she's out having lunch, or yknow people play games, you know that. And when somebody's........if there's a leak running down, a faucet leaking in the bathroom, and this individual knows that last time they tried to call and suddenly got their butt chewed, they're liable to turn their back and walk away from it, and saying no I'm not going to put up with that anymore. Chris takes just about all calls from everybody. Some people should be more in depth in gathering the knowledge about the call, then she might be willing to go to the building manager. We accept all the calls from all buildings in our zone.

RM: I guess the last question I have is just basically (about) this database that contains the information from these service requests. I'm wondering if it is at all possible for me to perhaps get access to that, so I can do some verification of data analysis on that for some of these subjective probabilities that we talked about today.

RH: Way over my head

RM: Ok

RH: That's....to be real honest with you , as much as we like supporting the students, postgraduates and stuff, in their work and stuff, that could be...that kind of information on that scale might not be well received by yknow a lot of people yknow.... I don't want.......This is over my head.

RM: I see

RH: Yea, in fact we're all passworded , and almost all the guys can get in and take a look, but then once you get into the point where you can enter information, it closes rapidly, and then once you get to the point where you can alter information already in there, that gets even smaller, and then, actually there are system managers that have unlimited access.

RM: Oh yea, so I mean as far as general searches, and data analysis & that sort of thing......

RH: Yknow I could pull up, I could tell you mainly like how many ....but see we...hot & cold fall under a lot of different things that we throw under one category, building ventilation. Ok so, I could pull that number and tell you we took 300 hits or something in the last month, but it wouldn't be a true direct reading of what you want, because we have more in there than just hot & cold calls.

RM: I see, because there's like too much air, etc....

RH: Well, the number we use is AC1089, and this is just called building operations & ventilation, Ok, but actually, that's kind of a generic catchall for the stationary engineers. OK, so it wouldn't reflect the true picture. You might want to submit a letter to John Rolle. He's the acting associate director of physical plant right now. Explain what you're doing,...... and ask him if you can meet with the FMIS (is what our system's called) manager, and the reason for meeting with him would be to gather background information, you know, for your postgraduate work. OK, and it would probably take somebody at John's level, if not higher for approval.