Proceedings of The ASME 2009 International Design Engineering Technical Conferences &
Computers and Information in Engineering Conference
IDETC/CIE 2009
August 30-September 2, 2009, San Diego, USA

DETC2009-87611

LOW-FREQUENCY PULSE WIDTH MODULATION DESIGN FOR HVAC COMPRESSORS

William Burke
PhD Candidate
Department of Mechanical Engineering
University of California
Berkeley, California 94720
Email: billstron@berkeley.edu

David Auslander [*]
Professor of the Graduate School
Department of Mechanical Engineering
University of California
Berkeley, California 94720
Email: dma@me.berkeley.edu

ABSTRACT

This paper discusses the design of low-frequency pulse width modulation for heating ventilation and air conditioning (HVAC) compressors. HVAC units are traditionally controlled using non-linear control techniques like hysteresis control. Using a very long pulse width, this method can treat an on/off air conditioner or heat pump compressor as a variable input for which traditional linear (or nonlinear) controls can be applied. The key advantage of this method is direct control over the compressor power using tunable saturation. Power control is especially useful when considering load management and real time energy pricing.

INTRODUCTION

The most common type of residential heating ventilation and air conditioning (HVAC) compressor is single speed, meaning it is either off or on at full power. Because of maintenance, reliability, and efficiency concerns, the compressors must cycle at relatively low frequencies. Furthermore, heat transfer dictates that the system, i.e. house, reacts slowly to the HVAC input and environmental inputs, meaning there is considerable residual in the system output (inside temperature). Traditionally, these systems use a non-linear hysteresis controller for temperature set-point following. The cycle rate is not directly defined, instead the width of the hysteresis band determines it. Hysteresis control is very simple to implement, model free, and robust. Unfortunately, it has a number of disadvantages when viewed from a modern perspective.

We are proposing a technique in which the single speed compressor can be treated as a variable power unit using low frequency pulse width modulation (PWM). Providing that the continuous system responds slowly, the discrete time PWM system can still be considered linear. The difficulty arises in error measurement because the states of the system could change considerably from the start of the PWM time period to the end. Consequently, the main design effort comes in appropriate filter design.

Low frequency PWM control has a number of advantages over traditional control of HVAC compressors. Firstly, any linear or non-linear control design technique producing a proportional input signal can be used to control the unit. Another advantage is that the power consumption of the unit can be explicitly controlled using tunable saturation limits, which is particularly important for load management and real time energy pricing. Finally, operation of multi-stage and variable HVAC compressors becomes much easier with a proportional control signal.

Very few people have written about low frequency PWM for residential HVAC compressors. A similar idea was presented in [1] for modulation of peak load in a multi-unit facility.

---

[*] Address all correspondence to this author.

MOTIVATION

The key motivation of this advancement came from the desire to simplify load management using thermostatically controlled devices. Firstly, load management is systematic modification of the load on the electricity generation and distribution system. Load management comes in many different flavors, but one key type is *peak shaving* which is needed when the power demand peaks above the generation capacity. A good overview of load management can be found in [2]. Furthermore, load management could (and will) be extended to provide marketable products like load following and spinning reserve.

Thermostatically controlled devices are well suited for load management because they are ubiquitous and heavy electricity consumers. Specifically, air conditioning contributes heavily to the summer peak loads through out the United States. Programmable Communicating Thermostats (PCTs) were recently proposed as a method to provide load management. PCTs were envisioned to be low cost residential thermostats with the ability to communicate with some central authority for the purpose of reducing power when needed. Using the model of the PCT, we want to extend their capability by providing more intelligence while keeping their price low.

Traditional non-linear control of thermostatically controlled devices complicates energy consumption analysis and manipulation. The inherent non-linearities make system identification and prediction difficult and unreliable. Furthermore, direct load control is crude and complicated with hysteresis control. Traditionally, it has been provided by communicating switches placed directly on the compressor to deny it power, bypassing the temperature controller altogether as in [3].

THEORETICAL BASIS

Pulse width modulation is a common technique for obtaining quasi-continuous output from an on/off type actuator. In its most common form, a PWM signal is a series of pulses produced on a fixed period (T). The on time (T_{on}) of the pulses varies between zero and full period. The varying pulse produces a variable output from the on/off actuator. If possible, a separate signal is used to control the direction of the actuator. PWM is generally specified in percent of period or duty ratio given by Equation 1.

$$\phi(kT) = \begin{cases} \frac{t_{on}}{T} & \text{for positive actuation} \\ -\frac{t_{on}}{T} & \text{for negative actuation} \end{cases} \tag{1}$$

The input to the system can be represented by the following:

$$u(t) = \begin{cases} U_{max} sgn(\phi) & \text{for } kT \leq t < kT + |\phi(kT)|T \\ 0 & \text{for } t \geq kT + |\phi(kT)|T \end{cases} \tag{2}$$

Now consider the linear time invariant continuous system represented by the standard state space formulation with n states, p inputs, and m outputs.

$$\begin{array}{ll} \frac{d}{dt}x = Ax(t) + Bu(t) & A \in \Re^{nxn} \;\; B \in \Re^{nxp} \\ y = Cx(t) + Du(t) & C \in \Re^{mxn} \;\; D \in \Re^{mxp} \end{array} \tag{3}$$

Let's say that one of the inputs to the system is given in terms of PWM. The system response to the discontinuous input $\phi(k)$ can be represented as in Equation 4.

$$x(t) = \begin{cases} e^{A(t-kT)}x(kT) + \int_{kT}^{t} e^{A(t-\tau)} BU_{max} sgn(\phi(k)) d\tau \\ \qquad \text{for } kT < t \leq kT + |\phi(k)|T \\ e^{A(t-kT-|\phi(k)|T)} x(kT - |\phi(k)|T) \\ \qquad \text{for } t > kT + |\phi(k)|T \end{cases} \tag{4}$$

By discretizing the system at the PWM sample time T, a single non-linear equation describes the response to the PWM input at the instances $t = kT$.

$$\begin{aligned} x((k+1)T) &= A_d x(kT) + h(kT, u) \\ A_d &= e^{AT} \\ h(kT,\ u) &= e^{AT}(I - e^{-AT|\phi(k)|}) A^{-1} BU_{max} sgn(\phi(k)) \end{aligned} \tag{5}$$

Linearizing the non-linear function $h(kT, u)$ yields Equation 6.

$$\begin{aligned} x(k+1) &= A_d x(k) + \hat{B}_d \phi(k) \\ A_d &= e^{AT} \\ \hat{B}_d &= (e^{AT} - I) A^{-1} BU_{max} \end{aligned} \tag{6}$$

Traditionally, Equation 6 is only considered a valid approximation when the PWM sample time (T) is small, but that is not entirely true. Actually, the matrix quantity AT must be small, giving rise to the possibility that the system matrix A is small and the sample time T is large. Note that similar analysis of PWM systems can be found in [4, 5] among others.

The resulting discrete time linear system is only so useful though. The value of the states and output at instant Tk, are just that, and the value of the states and output between $T(k-1)$ and Tk are not considered. However, given the long PWM period, the values between discrete instances are still important. This gives rise to a filter design problem. What is the best output or state filter considering controller performance and objectives?

SYSTEM DESIGN

For the design of the low frequency PWM controller, we will use two different house models. For the final design, we apply the controller, and do final tuning, on a relatively complicated house model used for load management experimentation. We previously outlined the model in [6]. For rough design, we consider the following first order continuous time multi-input single output system in state space form (Equation 7).

$$
\begin{aligned}
\dot{x}(t) &= Ax(t) + Bu(t) \\
y(t) &= Cx(t) + Du(t) \\
A &= [-4E-4] \quad B = [4E-4, -2.5E-6] \\
C &= [1] \qquad\quad\; D = [0,0] \\
u(t) &= [T_{out}, P_{ac}]^T \\
y(t) &= T_{in} \\
M &= 4000
\end{aligned}
\tag{7}
$$

The inputs to the system are the outside temperature and the instantaneous power from the single-speed HVAC compressor, $[T_{out}, P_{ac}]$. In this model, we approximate the compressor as producing full power ($M = 4000$) instantly. Further, the compressor produces the identical power regardless of outside temperature. The output of the system is the indoor temperature, T_{in}. The goal is for the output to track a set-point temperature, $y_{ref} = T_{sp}$.

For reliability, maintenance, and efficiency reasons, HVAC compressors should not be cycled too often, 4 to 6 times an hour. Considering this against the desires to have accurate inside temperature reference tracking and load management, we decided on a fifteen minute PWM period. Using this sample rate ($T = 900s$), we can discretize the continuous time system using Equations 5 and 6. Equation 8 describes the linear approximation of the system with the the variable u_{dr} representing the duty ratio of the HVAC compressor on the interval $[0,1]$. Without loss of generality, we have assumed that the first input, outside temperature (T_{out}), fluctuates slowly with respect to the sampling interval. (If the outside temperature were to fluctuate quickly, $\hat{B}_d(1,1)$ would be different and the signal would need to be appropriately filtered. However, no changes would be made to the PWM input or its effect on the system.) Figure 1 shows the simulation response to an open loop PWM signal for the continuous time system, sampled non-linear system, and sampled linear approximation.

$$
\begin{aligned}
x(k+1) &= A_d x(k) + \hat{B}_d \phi(k) \\
A_d &= [0.69768] \\
\hat{B}_d &= [0.30232, -7.5581] \\
\phi(k) &= [T_{out}, u_{dr}]^T
\end{aligned}
\tag{8}
$$

The previous analysis was simply to show that given a very small A matrix, the low frequency PWM system becomes approximately linear.

Low frequency PWM requires a couple of changes from its high frequency counterpart. In the normal case of high frequency PWM, the system acts as a low pass filter attenuating the high frequency changing input. In the case of low frequency PWM, there is a large residual that necessitates a synchronous filter on the feedback path that operates on the (approximately) continuous signal but is downsampled at the controller sample period. Note that filtering with a continuous time linear filter does not change the linearity of the PWM system as long as the filtered system matrix remains small. Furthermore, the PWM sample rate is so slow that in order to obtain good controller performance, the controller must run at the same rate as the PWM. With high frequency PWM, the PWM sample rate and controller sample rate can be chosen somewhat independently.

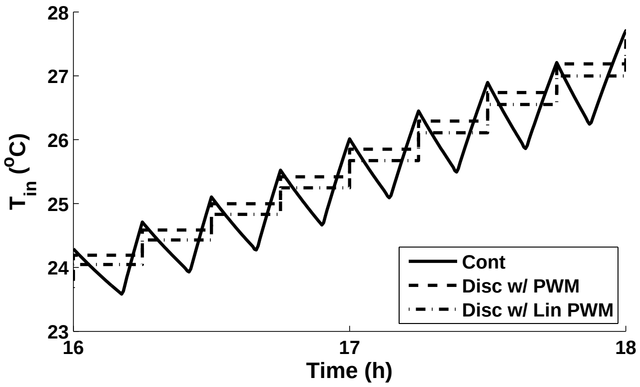

(a) Output

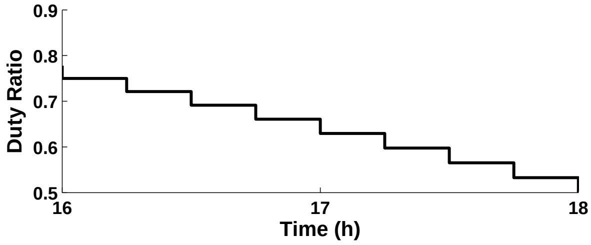

(b) Input

Figure 1. Discretization of Continuous System

Output Filter Design

It is certainly possible to design an analog filter to meet our requirements, but considering the long time scales, the electronic components would need to be very large and expensive. Therefore we will leave exact linearizability behind and restrict ourselves to approximation with digital filters. The filter sample rate (T_s) must be small compared to the PWM sample rate (T). Further, the filter should be synchronized with the PWM sample rate to ensure that the control receives the most up to date information. Hence, the filter sample rate must be chosen so that the PWM sample rate is an integer multiple of it.

Our goal for filtering is to attenuate the residual caused by the low frequency pulsed input. Explicitly, the synchronous filter should minimize the error over the previous time step (Equation 9). The synchronous filter signal is given by $y_f(kT)$, and the continuous time signal is given by $y(t)$.

$$\| e_f \| = \left(\sum_{k=0}^{\infty} \sum_{i=0}^{T/T_s-1} \left(y(Tk - T_s i) - y_f(Tk)\right)^2 \right)^{1/2} \tag{9}$$

Note that the output signal will be fluctuating in response to the PWM input at the the known PWM frequency. The PWM forced fluctuation is exactly the signal we want to attenuate. Therefore, we need a low pass filter with a cutoff frequency below the PWM frequency.

Table 1. Parametric Filter Study

Type	**n**	ω_n	$\| e_f \|$
butter	3	0.0007937	99.721
butter	3	0.0009259	92.55
butter	3	0.0011111	84.306
butter	3	0.0013889	75.853
butter	3	0.0018519	69.395
butter	3	0.0027778	67.069
butter	3	0.0055556	79.135
butter	1	0.0013889	65.648
butter	2	0.0013889	67.298
butter	4	0.0013889	87.528
butter	5	0.0013889	99.994
butter	6	0.0013889	111.29
boxcar	30	-	102.76
boxcar	60	-	84.633
boxcar	90	-	73.561
boxcar	120	-	68.748
boxcar	150	-	65.659
boxcar	180	-	**64.042**

There are two main classes of filters – infinite input response (IIR) and finite input response (FIR). The Butterworth filter is an example of a simple IIR filter. A discrete time Butterworth filter is designed using two parameters, cutoff frequency (ω_n) and system order (n). The Boxcar filter is an example of a simple FIR filter. We performed a parametric study of the Butterworh filter with varying parameters and the Boxcar filter with varying orders. Each of the filters is sampled at 5 second sample rate and downsampled to the PWM sample rate. Table 1 shows the norm error for each filter tested.

As expected, the choice of error norms played a large part in selection of the filter. The Butterworth filter is very sensitive to the choice of cutoff frequency, and it seemed to perform best at around $\omega_n = 0.0028 rad/s$. With a fixed cutoff frequency, the Butterworth filter performed worse as the order increased, mainly because of the increasing delay. The high order ($n = 180$) boxcar filter yielded the minimum norm error, and this is not surprising given the specification of the norm error. Another great advantage of this filter is ease of computation. It only requires storage of one variable.

Controller Design

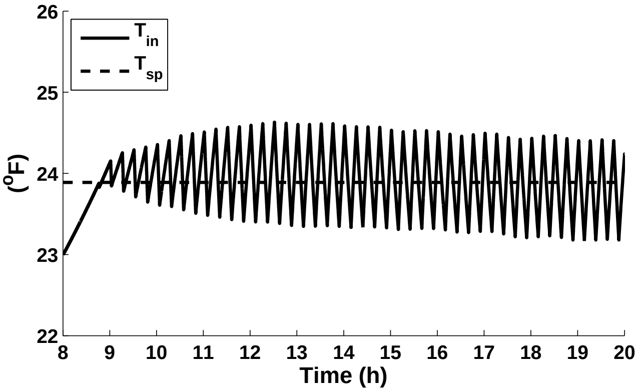

(a) Output

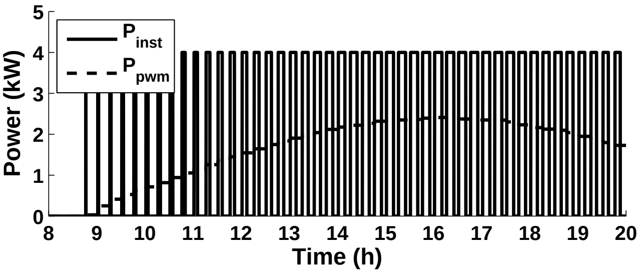

(b) Input

Figure 2. PWM Control of First Order System

The filtered system is approximately linear, and a controller can be designed using any linear (or non-linear) design technique. We chose to design a proportional plus integral (PI) controller. This type of controller is quite simple to design and program (Equation 10).

$$\begin{aligned} e(k) &= y_{ref}(k) - y_f(k) \\ e_{int}(k) &= e_{int}(k-1) + e(k)T \\ P(k) &= k_p e(k) + k_i e_{int}(k) \end{aligned} \quad (10)$$

The only trouble is how to deal with saturation. An HVAC compressor represents a single sided input, i.e. cooling or heating. Additionally, residential HVAC systems rarely control both the heating and the cooling systems at the same time. For instance, in cooling mode the temperature is allowed to drop well below the set-point. In these situations, the integrator in the PI controller winds up, creating poor performance. This necessitates an anti-windup mechanism, as illustrated with Equation 11.

$$e_{int}(k) = \begin{cases} \frac{P_{max} - k_p e(k)}{k_i} & P(k) > P_{max} \\ \frac{P_{min} - k_p e(k)}{k_i} & P(k) < P_{min} \end{cases} \quad (11)$$

In order to tune the controller, we used an iterative process on the first order system to get the order of magnitude of the gains. The final tuning was completed, iteratively, on the more complicated system. Figure 2 shows the results of a PI control on the first order system. In this and all of the first order simulation results, the outside temperature is time varying with a sine wave that peaks at 32.2^oC at 4:00pm and has a minimum value of 21.1^oC at 4:00am.

On/Off Time Limits

For maintenance and reliability reasons, typical HVAC compressors need to be in the on state or off state for a certain amount of time before switching. This requirement is typically accomplished using on/off timers directly on the compressor unit. These traditional cycling timers will still work with low frequency PWM actuated units, but the control will be slightly biased as a result. Further, low frequency PWM can account for these timers directly by making use of slightly more complicated saturation guidelines that round the low and high PWM to ensure the on/off times are met.

Multi Stage Units

Multi-stage compressors have been commonplace for years. Without going into the mechanical design of these units, multi-stage compressors essentially have more than one output power that can be switched between. In general, they offer pretty significant efficiency advantages over traditional single-stage units.

Unfortunately, operation of multi-stage units using tradition hysteresis control is cumbersome at best. Using hysteresis control, the controller has no way of judging how much power is needed to control the system. Some additional rate detection, etc. needs to be implemented in order to make hysteresis control feasible.

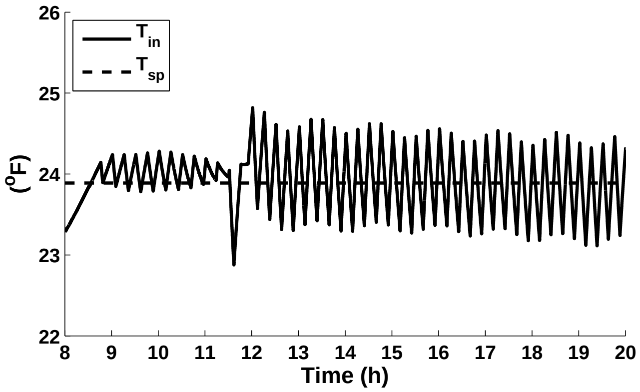

(a) Output

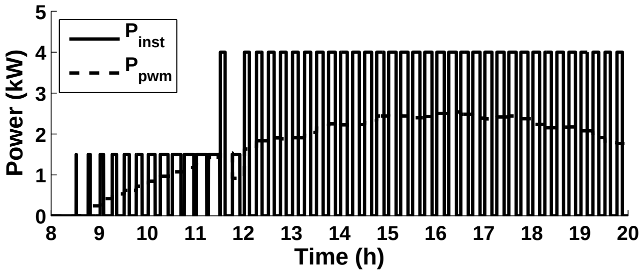

(b) Input

Figure 3. PWM Control with 2 Stage Compressor

Control of a multi-stage unit using low frequency PWM is very simply accomplished. When creating the PWM duty ratio from the control calculation, $P(k+1)$, the largest stage with power less than $P(k+1)$ should be used. Equation 12 illustrates the calculation for a two stage unit. Figure 3 shows simulation of the same first order system as in Figure 2, with the identical controller (gains included), but with a two stage compressor.

$$\begin{aligned} P_{cur}(k+1) &= \begin{cases} P_1 & 0 < P(k+1) \le P_1 \\ P_2 & P_1 < P(k+1) \le P_2 \end{cases} \\ u_{dr}(k+1) &= P(k+1)/P_{cur}(k+1) \end{aligned} \quad (12)$$

Tunable Saturation

Control of HVAC power consumption usually takes the form of radio operated direct load control (DLC) switches attached to the compressor, as in [3]. A DLC switch bypasses the temperature controller and shuts off the compressor for a specified interval of time. In general, the switch does not consider the cycling characteristics of the unit and simply shuts off when commanded. This results in dramatically different responses for different compressors, ranging from almost no change at all if the

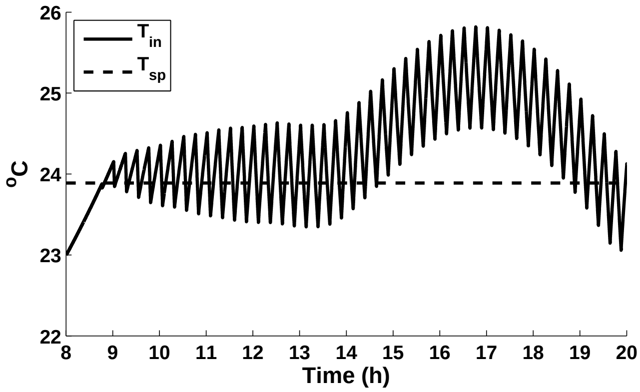

(a) Output

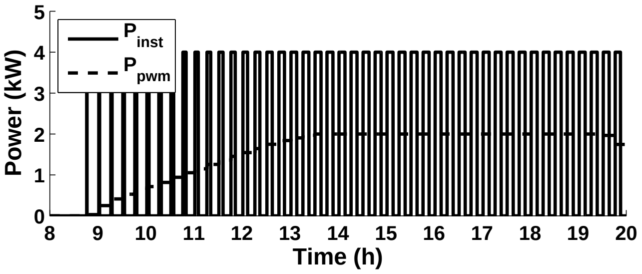

(b) Input

Figure 4. PWM Control with Tunable Saturation

natural cycling period is below the switch setting, to dramatic changes when the natural cycling period is much greater. Some DLC manufactures have tried to remedy this problem by introducing adaptive switches that reduce the power proportionally, as in [7].

One of the key advantages of PWM actuation is direct control over power consumption using tunable saturation. Figure 4 illustrates this concept using the same first order system as used in the previous examples. Here, the saturation level was statically set at 50%, but it would be straightforward to extend the system with a higher level controller manipulating the saturation in real time.

RESULTS

In order to verify the performance of the PWM actuated system, we performed experiments using a more complicated model that better simulates the dynamics of a house. We previously described this model in [6]. Two tests were performed on an identical house model under identical environmental conditions, i.e. outside temperature and solar radiation from a hot summer day in Fresno California. The first test controlled the HVAC using a hysteresis control, and the second test used the PI controller with PWM actuation as designed previously. Figure 5 illustrates the results of the two controllers tracking an identical set-point temperature.

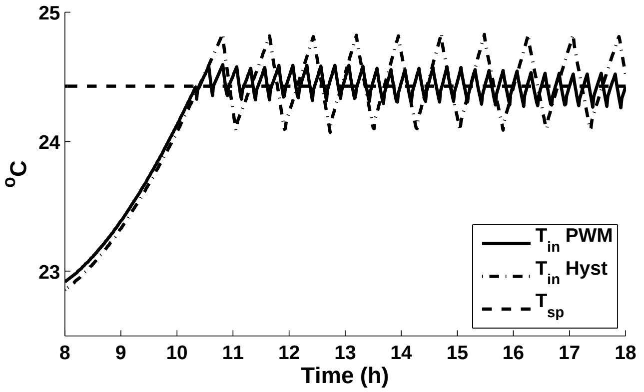

(a) Output

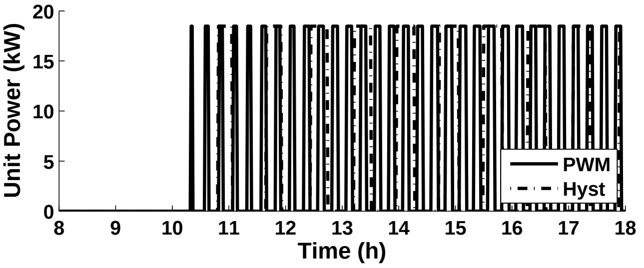

(b) Input

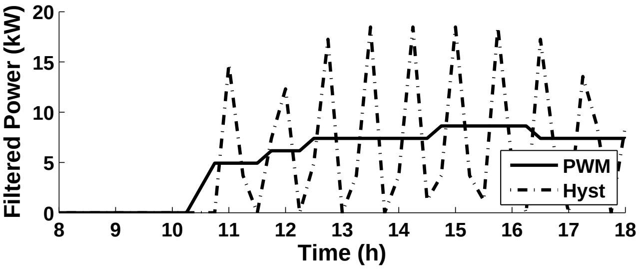

(c) Filtered Input

Figure 5. PWM Control of Full Simulation

The first obvious difference between the two controller types is the difference in error between the on and off peaks – the error band. For the PWM actuation, the error band is “set” by the PI-controller *and* the choice of PWM frequency. If it were hotter outside, causing the temperature in the house to rise more quickly, the unit would cycle at the same rate because the PWM frequency is fixed, but the error band would be larger. At lower outside temperatures, the frequency would still be the same and the band would be smaller. Alternatively, the error band for the hysteresis controller is set *directly* by the controller *only*. Regardless of the outside temperature, the error band will always be the same (if the compressor has the capacity to cool the house), but the cycling frequency fluctuates. If it is hotter outside, the unit cycles more quickly, and cooler temperatures result in less frequent cycling. This fluctuating cycling rate makes prediction and analysis difficult because of the lack of time consistency.

The linearizing quality is the main advantage of PWM actuation. Figure 5(c) plots the filtered power using a boxcar filter over a fifteen minute interval similar to the one used in the controller. This treatment of the system input clearly shows the discontinuous hysteresis control and the smooth PI control.

CONCLUSION

We demonstrated, through simulation, how low frequency PWM simplifies control of multi-stage compressors. The compressor stage is stepped based on a simple set of PWM rules. The simplicity advantage extends for variable-speed compressors a well. It is in fact simpler as the compressor speed is determined directly via the controller.

Low frequency PWM control dramatically simplifies the analysis and control of HVAC compressors when viewed through the lens of load management. Hysteresis control results in a difficult to analyse highly non-linear system. System identification is difficult, meaning that state prediction is unreliable. PWM control linearizes the system, simplifying not only controller analysis but system identification and prediction as well. Our future research will take advantage of these key properties of the PWM actuated system.

Assume that energy consumption is roughly proportional to the compressor on-time. This is mostly true except that the efficiency (and therefore power) varies somewhat with outdoor temperature. With low frequency PWM, the control signal is calculated at the start of the PWM period, and therefore the energy consumption for the period is know in advance. This results in the ability to artificially limit the power consumption using a simple tunable saturation variable. This is a major advantage for PWM actuation that, when coupled with the linearizing qualities, will enable more intelligent load management systems than could be designed using hysteresis control. Our further research is aimed exactly at this.

ACKNOWLEDGMENT

Thanks go to Ron Hoffmann and Joe Eto for their excellent advice and encouragement.

This work was supported by the California Energy Commission, California Institute for Energy and the Environment under project number C08019. It does not necessarily represent the views of the Energy Commission, its employees, or the State of California. The Energy Commission, the State of California, its employees, contractors, and subcontractors make no warranty, express or implied, and assume no legal liability for the information in this paper; nor does any party represent that the use of this information will not infringe upon privately owned rights.

REFERENCES

[1] Federspiel, C., Lanning, S. D., Li, H., and Auslander, D. M., 2001. “Coordinated control of HVAC systems”. In ICEBO - International Conference for Enhanced Building Operations, Energy Systems Laboratory (http://esl.tamu.edu).

[2] Bellarmine, G. T., 2000. “Load management techniques”. In IEEE SoutheastCon 2000 ’Preparing for the New Millennium’, Apr 7-Apr 9 2000, Conference Proceedings - IEEE SOUTHEASTCON, Institute of Electrical and Electronics Engineers Inc., Piscataway, NJ, USA, pp. 139–145.

[3] Navid-Azarbaijani, N., and Banakar, M., 1996. “Realizing load reduction functions by aperiodic switching of load groups”. *IEEE Transaction on Power Systems, Vol 11, No. 2*.

[4] Sira-Ramirez, H., 1989. “On the design of nonlinear pulse-width-modulation controllers”. In Control and Applications, 1989. Proceedings. ICCON ’89. IEEE International Conference on, pp. 397–402.

[5] Hou, L., and Michel, A. N., 2001. “Stability analysis of pulse-width-modulated feedback systems”. *Automatica,* **37**(9), Sept., pp. 1335–1349.

[6] Burke, W. J., and Auslander, D. M., 2008. Modular and extensible systemic simulation of demand response networks, October.

[7] Moe, M. L., and Hedman, T. A., 1997. Adaptive load cycler for controlled reduction of energy use, Oct. U.S. Classification: 364/492; 364/483; 364/557; 307/32; 307/39 ; International Classification: H02J 318; G05F 166.