Using Demand Based Reset Strategies

Reinhard Seidl, P.E.
Taylor Engineering

Abstract

Control systems that rely on fixed setpoints to achieve stable operation typically result in higher energy costs than systems operating with variable setpoints. Typical examples of such setpoints are supply temperature for an air handler, supply pressure for an air handler, and chilled water temperature for a chiller plant.

Fixed setpoints are sometimes used on purpose to achieve stability, simplicity and reliability at the expense of efficiency. This may be the case for critical facilities such as labs or clean rooms. There are also cases where the method of using variable setpoints is not implemented because it was not known to the designer, or because, more commonly, it is implemented incorrectly and turns out to be ineffective.

This paper addresses the basic methods of resetting setpoints to match the actual building or zone demand, and our experience in commissioning systems with demand-based resets to ensure that they work.

About the Author

Reinhard Seidl is a Principal at Taylor Engineering, an engineering consulting firm based in Alameda, California. His work over the last 2-3 years has focused to a large degree on control systems and commissioning of buildings. In this role, he has been an active participant in the development of the Universal Translator (www.utonline.org and http://tedownloads.com/UT/training.shtml), a free publicly available tool developed in collaboration with PG&E to enable building owners, engineers, and commissioning providers to more effectively use information produced by building energy management and control systems (EMCS). He is also a member of ASHRAE's guideline committee GPC1.2, at present developing a guideline for commissioning existing buildings, and ASHRAE's technical committee 7.5 on smart building systems.

Introduction

Thanks to the growing application of energy management and control systems (EMCS), it has been become possible to monitor equipment much more closely than was possible with pneumatic and electric control systems. When an EMCS includes controls "at the zone level", we typically mean that controls extend beyond the central mechanical rooms to individual air terminals serving spaces in all parts of the building. For new installations, this is usually the case.

For retrofits or upgrades of existing facilities however, extending an EMCS to the zone level is often too costly, and only individual pieces of equipment or mechanical rooms are equipped with modern controls. This type of retrofit for older systems does not allow use of demand based reset as described in this paper.

We will clarify the methods used to implement demand based reset strategies where the EMCS infrastructure allows, and discuss how to make these systems work as intended.

Demand Based Reset Strategies

Background

A traditional approach to controlling systems is to specify fixed setpoints. A typical example would be to run a rooftop package unit with a fixed supply air temperature setpoint of 55°F and a fixed supply pressure setpoint of 1.5" wg. This assures stable operation and reduces comfort problems and callbacks. It also assures less than optimal energy efficiency.

Since 1999, ASHRAE Standard 90.1 requires the use of a demand-based static pressure setpoint reset for air systems with zone level controls, as does California's energy code, part 6 of Title 24. The table below shows a quick overview of the regulatory landscape.

Table 1: Regulatory requirements for air systems

System type	Air	
Reset mechanism	Pressure setpoint reset	Temperature setpoint reset
Standard 90.1-2004[1]	Required, demand-based, §6.5.3.2.3	Not required
2006 IECC[2]	Required, demand-based, §503.4.2	Not required
2005 California T24[3]	Required, demand-based, §144.(c).2.D	Required except for systems with variable speed drives, §144.(f)

Table 2: Regulatory requirements for hydronic systems

System type	Hydronic	
Reset mechanism	Pressure setpoint reset	Temperature setpoint reset
Standard 90.1-2004	Not required	Required except for variable flow systems, §6.5.4.3
2006 IECC	Not required	Required except for variable flow systems, §503.4.3.4
2005 California T24	Not required	Required, except for variable flow systems, §144.(j).4

As can be seen, air systems are typically required to have pressure setpoint reset while not having temperature setpoint reset. Hydronic systems are the reverse: regulations require the use of temperature setpoint reset while not requiring the use of pressure setpoint reset.

For an explanation of this, we'll begin by looking at indicators used for building demand, and the resulting reset mechanisms that make sense.

Building Demand

To begin with, let's describe how we can quantify building demand, because that's what we have to respond to. The following list describes typical demand indicators:

Poor Demand Indicators for Supply Air Pressure Reset

- Air terminal damper end-switch (in the absence of a flow sensor, on legacy systems): this only indicates that a damper is wide open. The damper may at that point be receiving in excess of the design airflow rate, so this is not a good indicator of required pressure.
- Approximate damper position: on floating actuators, which receive only "open" or "close" commands, there is no direct indication of damper position. By summarizing the "open" and "close" commands given over time, the damper position can be approximated but is usually not very accurate, if the damper is not driven to an auto-zero position at regular intervals (example: once every 24 hours).

Good Demand Indicators for Supply Air Pressure Reset

- Comparison of airflow to setpoint: in lieu of approximating damper position, a call for pressure can be generated when the airflow of a zone is below setpoint for a certain period of time. This method can be used when damper position is unavailable. Note that this indicator can only be used for the trim-and-respond logic described below, not for the PID[4]-based reset logic.
- Damper position: when pressure-independent terminals (with flow sensor) are used, a fully open damper is an indication of insufficient flow. Note this is not the case in terminals without flow sensors, see above. Damper position can be sensed by using a modulating actuator, or by using a floating point actuator with position feedback.
- Floating point actuator position, approximated by summarizing "open" and "close" commands over time, and accompanied by driving the damper to zero position for several seconds at regular intervals to eliminate inaccuracy in the approximated position.

Poor Demand Indicators for Supply Air Temperature Reset

- Return air temperature: Provides an average of what is occurring throughout the entire building, or at least a wing, floor or large section of the building. As such, we will easily overlook a number of rooms that are not receiving sufficient cooling or heating, because their operation is lost in the averaging effect.
- Typical zone sensors: This is a technique often used for retrofits of older pneumatic systems when installing zone-level EMCS is too expensive. One sensor may be installed at the top floor in a North-facing zone with the intent to reflect the coldest area in the building. Another sensor may be installed in the southwest corner of the building, with the intent to reflect the warmest area in the building. The system then judges whether the building is "too cold" or "too hot" based on the readings of these typical sensors. The sensors may become bad indicators due to churn in the area in question, or due to local load changes like small space heaters that provide local comfort.
- Outside air temperature: For buildings with large internal or solar load components, this may lead to incorrect adjustments. For example, take a building with high internal loads

that are constant throughout the day. Re-setting supply air temperatur (for energy savings through less cooling) at a cooler time of day may not work: thanks to the high internal loads, we may end up overheating the spaces served. Since we're only looking at outside air, and not the spaces themselves, this is a poor overall indicator of building demand.

Good Demand Indicators for Supply Air Temperature Reset

- When an air terminal in cooling meets the required airflow but fails to achieve comfort conditions in the space served, it requires colder supply air. The same is true in reverse for heating. Thus, the output of the cooling or heating PID loop of an air terminal serves as a good demand indicator for supply air temperature.

Poor Demand Indicators for Chilled/Hot Water Pressure Reset

- Valve position is generally a poor indicator of pressure demand, since no flow measurement is available at most control valves for coils. Thus, a wide open valve may receive in excess of design flow and would not be a good indicator for additional pressure.

Good Demand Indicators for Chilled/Hot water Pressure Reset

- In a hydronic campus heating system, each building may have a heat exchanger with a Btu meter indicating flow. If this is used, the measured flow can be compared to design flow and used to indicate demand for more pressure.

Poor Demand Indicators for Chilled/Hot Water Temperature Reset

- Return water temperature provides an average of multiple coils and will not reflect insufficient heating or cooling of particular coils.
- Typical zone sensors (see supply air temperature reset)
- Outside air temperature (see supply air temperature reset)
- Output from the heating or cooling PID loop: unlike the good indicator for air temperature reset, in hydronic systems we typically lack flow sensors. Thus, a heating or cooling loop that is not satisfied may be caused by insufficient flow. As such, neither the PID loop output nor the valve position are very good indicators for reset strategies.

It now becomes clear why regulations focus on pressure based reset in air systems while not requiring pressure resets in hydronic systems: there are good indicators for pressure demand in air systems, but not in hydronic systems.

Similarly, for temperature resets, there are no good indicators in hydronic systems. There are some poor indicators, and they may be used, but do not have to be used if variable flow systems are employed. The intent behind this language is that the majority of energy savings will come from use of variable flow which may be attained by either fixed pressure setpoint or variable pressure setpoint.

For air systems, there is a good indicator of demand, however, this is not mandated in any of the regulations listed in Table 1. This is because a temperature reset range has to be chosen well to ensure energy savings. Providing very warm supply temperatures has the potential to undo cooling energy savings through increased fan energy. However, this is not inherent in the principle of the reset, and depends on the selected setpoints rather than the method itself. We will return to this item in more detail.

Sequences of Operation

Using the indicators of building demand defined above, we can now create sequences of operation that react in accordance with changing building demand. There are two sets of sequences listed below. One addresses pressure reset, and one addresses temperature reset.
Each set of sequences is written in two ways, using either a PID-based reset method or a trim-and-respond based reset method. We will look at both methods in detail.

PID Based Reset Control Sequences

- Supply Air Pressure Reset: The static pressure setpoint shall be reset by mapping the output of a direct acting PID loop to a setpoint range of 0.15 inches (0% PID output) to 1.5 inches (100% PID output). The input of the PID loop is the most open damper with a setpoint of 90% (adjustable). When the fan is proven off, disable the PID loop to prevent wind-up and freeze the output at zero. The loop shall be tuned to react slowly.

- Supply Air Temperature Reset: During occupied mode, the supply air temperature setpoint is reset from T-min (53°F) when the outdoor air temperature is 70°F and above, proportionally up to T-max when the outdoor air temperature is 60°F and below. T-max shall be reset by mapping the output of a direct acting PID loop to a setpoint range between 65°F (0% PID output) and 55°F (100% PID output). The input to the reset loop is the highest zone cooling loop output with a setpoint of 99% or greater (adjustable). When the fan is proven off, disable the PID loop to prevent wind-up and freeze the output at zero. The loop shall be tuned to react slowly.

Trim-and-Respond Based Reset Control Sequences

- Supply Air Pressure Reset: Static pressure setpoint shall be reset using trim and respond logic within the range 0.15 inches to 1.5 inches. When the fan is off, freeze setpoint at the minimum value (0.15 inches). While fan is proven on, every 2 minutes, decrease the setpoint by 0.04 inches if there is one (adjustable) or fewer pressure requests. If there is more than one (adjustable) pressure request, increase the setpoint by 0.04. Where VAV zone damper position is known, a pressure request is generated when any VAV damper served by the system is wide open until it drops to 90% open. Where VAV zone damper position is unknown, a pressure request is made when the ratio of the zone's actual supply airflow to supply airflow setpoint is less than 90% until it rises to 100%.

- Supply Air Temperature Reset: During occupied mode, the supply air temperature setpoint is reset from T-min (53°F) when the outdoor air temperature is 70°F and above, proportionally up to T-max when the outdoor air temperature is 60°F and below. T-max shall be reset using trim and respond logic within the range 55°F to 65°F. When fan is off, freeze T-max at the maximum value (65ºF). While fan is proven on, every 2 minutes, increase the setpoint by 0.2ºF if there is one (adjustable) or fewer zone cooling requests. If there is more than one (adjustable) cooling request, decrease the setpoint by

0.3ºF. A cooling request is generated when the cooling loop of any zone served by the system is greater than 99% until it falls to 90%.

Practical Experiences

Loop stability

The sequences of operations listed above will all work. We have found through practical experience however, that the trim-and-respond logic offers some advantages.

For one thing, it is much easier to tune: the reset loops have to act very slowly in order to function correctly. Quite often, an installer or controls technician in the field will install PID-based logic using the default "out of the box" PID parameters, which are typically much too fast for a reset loop. As a result, the loop will reset every few minutes, as if it were a reheat valve responding to VAV discharge temperature. This kind of response speed is much too fast and will make the entire building system "oscillate" as all components become unstable.

The reset loop should normally work in cycles of hours, or even in a daily cycle. After all, its ultimate function is to adjust for the building demand as it varies due to weather or internal loads. An occasional peak may be present due to a particular load shift, but the reaction will typically be slow. Tuning a PID loop to respond in a matter of hours rather than minutes requires that the proportional gain (P) and integral gain (I) be set to values that are between 10 and 100 times lower than typical PID loop values. Finding the correct balance of PID parameters that provide stable operation and sufficient response speed can take several days to weeks of tuning, and requires a trained operator.

The figures below show an example of a chilled water loop with demand-based pressure reset using valve position, aimed at keeping the most open chilled water valve at 90% setpoint. As mentioned in the section on indicators, using valve position is a poor indicator, but it still works slightly better than a fixed pressure setpoint as long as there are no "rogue" zones (due to fouling or design problems).

Figure 1: Unstable demand-based pressure reset loop

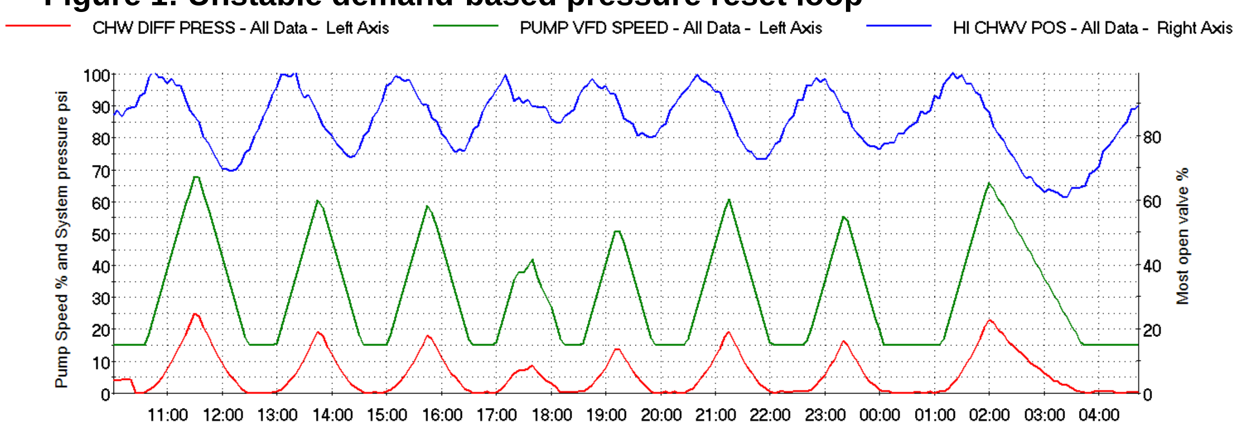

Figure 2: Stable demand-based pressure reset loop

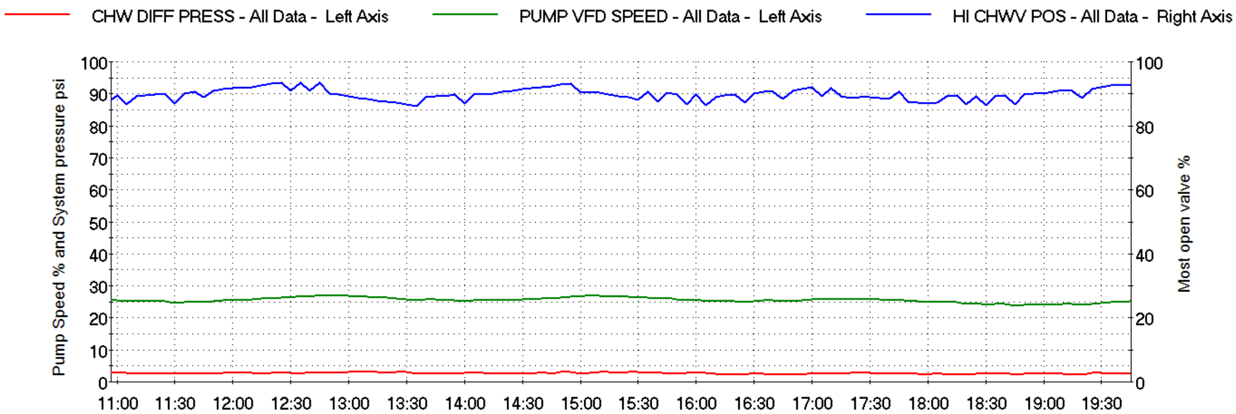

This loop required about a month to tune correctly, but ended up cutting pump energy by about 85%, as the main pump ultimately ran at an average of about 25% speed, down from the initial (fixed) speed of 50% that was selected before a demand-based loop was implemented. (Note that, in this case, 3-way valves were also converted to 2-way valves, so that the cubic equation of power vs. pump speed held true, as opposed to systems where only pressure is changed but not flow, in which case the correlation is linear).

During this month of tuning, only intermittent access to the system was required for minutes at a time to examine results from previous days or to look at performance changes over the weekend vs. weekdays, and to change PID parameters slightly. For this kind of work, a web interface is vital, since it can limit the amount of time required for such adjustments to literally 30 minutes, whereas visiting the jobsite might take half a day or more.

In general, a trim-and-respond loop will not require as much tuning as a PID-based reset method, and in addition, the reaction speed can be set to fast for ramping up pump or fan speed (example: to react to Monday morning occupancy) while it can be set to slow for ramping down (to help stabilize a loop). In general, a PID loop does not allow tuning of this kind, as the reaction speed is inherently linked to the process signal and the set of fixed PID parameters.

Simultaneous temperature and pressure reset

The question often arises whether it makes sense to have a temperature reset running simultaneously with a pressure reset. After all, by raising supply temperature in an air system, more air has to be supplied for the same cooling effect. So doesn't this cancel out the savings? And how will the loops keep from "fighting" each other?

Let's begin by looking at the issue of the two reset mechanisms "fighting" each other. This only occurs as a result of an unrelated issue, namely when loops are badly tuned. Even then, the instability that is witnessed is really a result of tuning and not of loop interaction. The incorrect perception is the following: by raising supply air temperature in the temperature reset loop, we force air terminals to request additional pressure, thus forcing the pressure reset loop to raise pressure. As the pressure reset reaches its limit, the zones still need more cooling, and now the temperature loop has to reset back down. As soon as it does, the zones are overcooled, dampers close, and the pressure reset loop is forced to lower static pressure in turn. And so on. As we just mentioned, this effect only occurs when loops are tuned badly. If supply air temperature is reset up when flows are met through the pressure reset loop, this will not make the pressure loop unstable. It may increase supply pressure somewhat, but it is more likely that there is one zone that is driving the pressure loop while a completely different zone is driving the temperature loop, and so the two are not related.

The question regarding the energy efficiency of running a temperature reset loop is more complex:
As an example, imagine an air handler that uses supply air temperature reset. It supplies 10,000 cfm of 65°F air to spaces that are kept at 75°F setpoint. The heat removed from the building is therefore roughly 10,000 * (75°-65°) * 1.08 = 108kBtu/h or 9 tons of sensible load. The outside air temperature is 75°F, and chilled water is provided by the central plant at 0.5 kW/ton. With mixed air temperature at 75°, the unit consumes roughly 10,000 * (75°-65°) * 1.08 = ~~140kBtu~~108kBtu/h or 9 tons of cooling (ignoring latent heat), for 4.5 kW of energy used[e1].[Rev-1-2] The fan is running at about 0.8 W per cfm. This means the unit uses 8 kW of fan energy and 4.5 kW of chiller plant energy for a total of 12.5 kW.

If there were no temperature reset, and the unit was providing 55°F supply air, we would need only half as much supply air since we have doubled the ΔT between supply air temperature and room temperature. This would result in fan power of $1/2^3$ or 1/8 of the original, or 1 kW if we are using perfect fan laws in an ideal duct system. We now need to cool the supply air by about 18 tons (again ignoring latent heat). We are thus using 9 kW of cooling energy and 1 kW of fan energy, for a total of 10 kW energy consumption at the unit. This is a slightly better energy picture than when temperature reset is used. However, this only holds true when we ignore reheat as a factor in this energy equation.

The real picture is more likely as follows: in a typical building, some of the low-load zones (North facing) were already at less than full airflow (for example, 40% of design flow) when the unit was at 65°F supply air. As the supply air temperature drops, these zones now drop to their heating minimums (for example, 30% of design flow) and begin to reheat to prevent being overcooled. Thus, the unit does not actually turn down to half of the supply volume, and we have to add reheat energy to the fan and compressor energy. There will also be more latent energy used for dehumidification, a factor previously ignored that depends on outside air humidity and will play less in dry climates than in wet climates. Thus, the actual system energy consumption of the unit might be 2 kW for fan energy, 10 kW for cooling energy, and 2 kW for reheat energy, for a total that is slightly higher than the energy used with full supply temperature reset.

The question becomes what amount of reheat will take place, whether this is done by hydronic heating or electrical heating, and what utility rates apply. The higher the supply temperature is reset, the higher the fan energy component will become, and the temperature reset limits that are chosen should reflect the balance between the energy consumed in different parts of the system.

Clearly, this is a much more complex mechanism than the pressure reset that is mandated by code. It is for this reason that supply air temperature reset was removed from Standard 90.1 as a requirement. Nevertheless, it can serve to reduce energy use when correctly set up and tuned, and becomes more effective in systems that carry a large reheat penalty, such as VAV reheat systems. It is less effective in systems with a lower reheat penalty such as fan-powered zones or dual duct systems, and is not recommended in systems without a reheat penalty such as single-zone applications.

Package units (as opposed to chilled water air handling units) generally do not lend themselves to the use of supply air temperature reset – regardless of unit capacity, the minimum cooling stage (typically one compressor, or at best, one compressor with hot-gas bypass control) is large enough to drastically change supply air temperature. As an example, a typical 50 ton unit with (2) 15 ton compressors and (2) 10 ton compressors operates at part load. It engages the first cooling stage, equivalent to 10 tons, in an attempt to reset supply air temperature from 63°F to 61°F. The unit design capacity is meant to bring 100% of rated airflow from 80°F to 55°F, or 25°F ΔT. At 50% of design airflow where the unit operates for our example, the 10 ton compressor capacity instantly reduces supply air temperature by 10°F (cooling cap/airflow = ΔT). So instead of resetting from 63°F to 61°F, the unit resets to 53°F, and will typically remain at that temperature for at least five minutes or so to prevent compressor cycling. Depending on the compressor staging, airflow and actual loads, these temperature swings can be far more drastic. In our experience, attempting temperature based reset with packaged units is not a successful strategy. The advent of digital scroll compressors and small variable speed centrifugal compressors may change this picture in the near future.

Rogue Zones

One problem frequently encountered with demand based reset loops is the so-called "rogue" zone. Imagine an IT closet designed to operate at 75°F. Tenants have set the thermostat to its lowest setting (60°F) for fear that equipment might be damaged at high temperatures. As a result, the room controls try to maintain 68°F which is substantially below design, and thus the zone generates cooling demand at all times. If this zone is recognized as a valid "player" in the building demand strategy, the air handler will never reset supply temperature, and all projected energy savings will be null and void.

Another type of rogue zone is the zone that is undersized by design. It will not attain design airflow unless pressure is raised to the maximum value in the reset range (and maybe not even then) and thus prevents the air handler from ever reducing pressure.

To avoid this scenario, we have to be able to ignore one or more zones – the trim-and respond sequences of operation shown earlier make allowance for this by allowing the number of zones that result in a system response to be adjusted by the user.

An even better approach in dealing with rogue zones is to make the EMCS programming flexible enough to identify the actual individual zones causing problems, and to allow the operator to either fix them or take them out of the building demand process altogether. This is unfortunately a very complex process to implement for most EMCS since the variables used in the EMCS equations have to be altered through the normal operator interface. Most control systems lack the ability to make such programming changes easy enough for an operator.

A simple approach to identifying a rogue zone is to program low level alarms into the EMCS – all systems have the capability of setting up alarms quite easily. To alarm for pressure-based rogue zones, simply set up alarms that are triggered when a zone damper or a valve opens to more than 95%. To alarm for temperature-based rogue zones, alarm whenever a cooling or heating loop reaches 95% or more. During initial commissioning, the alarm summary page may display just a few zones that alarm very repeatedly. If this is the case, these zones should be examined and if they cannot be corrected, they should be taken out of the reset sequence while the programmer is still on site.

For normal operation, the alarms should be intermittent in nature (a few per day) and should show a mix of different zones. Once this type of operation is attained after a few weeks, these alarms can be turned off. They can be re-enabled later if an air handler or pump is found "stuck" at maximum pressure setpoint, or at minimum temperature setpoint for periods of a day or more.

Summary

Demand based reset strategies offer the possibility of substantial energy savings when compared to operating equipment with fixed setpoints. To use these strategies correctly, a good set of building demand indicators has to be chosen, and the reset mechanisms have to be carefully programmed and commissioned in order to work.

In particular, they have to be tuned to remain stable – unlike many other building control mechanisms, they are rather unforgiving when it comes to stability and have to be set up carefully and individually. In addition, rogue zones have to be identified and eliminated to allow resets to work – depending on the system setup, a single rogue zone can completely eliminate the benefit of using demand based resets. This paper describes how to implement demand based reset mechanisms.

[1] ASHRAE Standard 90.1-2004, "Energy Conservation in New Buildings except Low-rise Residential Buildings." American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. Atlanta, GA, see

[2] 2006 International Energy Conservation Code, International Code Council, see http://www.iccsafe.org/e/prodshow.html?prodid=3800S06

[3] Title 24 part 6, 2005 Building Energy Efficiency Standards, California Energy Commission, see http://www.energy.ca.gov/title24/2005standards/index.html

[4] See PID basics at 2005 ASHRAE Fundamentals chapter 15