Energy Efficiency and the Control and Simulation of VAV Systems

Paul Bannister, Exergy Australia Pty Ltd

ABSTRACT

Variable air volume (VAV) systems are one of the most common air-conditioning system types for larger buildings in warm temperate climates and are likely to remain so for some time. However, VAV systems are heavily dependent upon control for their efficient operation and are particularly prone to system-wide failure as a result of the malfunction of individual components in the field.

In this paper the key control methodologies for VAV systems are discussed from the perspective of energy efficiency, robustness to common mechanical failures and long term manageability in a realistic operating environment. Key issues examined include fan control, supply air temperature control, VAV terminal control and the coordination of terminal and AHU actions to minimise simultaneous heating and cooling.

Unfortunately, many of the key algorithms used in VAV system control cannot be properly modelled in leading simulation packages, leaving most control optimisation to be assessed on the basis of intuition rather than analysis. This paper concludes with a list of key features of VAV control that should be included within simulation packages to enable the complete and accurate assessment of this important air conditioning system type.

Introduction

Variable air volume (VAV) air conditioning is one of the most common air-conditioning types of air-conditioning in warm temperate climates. This commanding position in the market has arisen over the past 15 years due to perceived comfort, flexibility and cost benefits, compounded by advances in building management and control system technology.

A further benefit perceived for VAV systems is that of energy efficiency, as all major components are equipped with the ability to turn down or otherwise adjust in response to load variations. A contrary view to this, however, is to consider VAV systems as being essentially little more than a series of modifications to what is fundamentally a constant volume air-conditioning system. In reality, both viewpoints are justified to at least some extent – VAV systems have the potential to be efficient but in practice this efficiency rarely appears to be fully realised, for reasons discussed in the balance of this paper. The challenge, therefore, is to develop the design and control of VAV systems in such a manner that the efficiency potential can be realised on a repeatable basis.

This paper summarises a range of approaches that have been applied in warm temperate Australia (which has a climate reasonably similar to coastal regions of California) with a view to improving the control and controllability of VAV buildings in order to achieve enhanced energy efficiency in operation.

Context

The office building sector in Australia has been significantly transformed over the past decade through the introduction of the Australian Building Greenhouse Rating (ABGR) scheme [1], which assesses the greenhouse performance (and consequentially the energy efficiency) of buildings in operation based on actual energy use, floor area, hours of operation and climate zone. This scheme is similar to the US EPA EnergyStar scheme for commercial buildings [2], except that in Australia, the common metering arrangements for buildings allow the "base building" (being the air-conditioning, lifts and common area services) to be separated from the "tenants" (being the lighting and plug-in power consumption of the tenants). By having separate ratings for landlord and tenant, a degree of market transformation has occurred such that there is now significant demand by tenants for rental accommodation with high base building ratings [3].

This process has highlighted the range of efficiency present across apparently similar buildings. Modern, apparently well designed VAV office towers have been shown to perform across almost the full 5 star range of the ABGR scale, from 0 stars to 4.5 stars [4]. In Sydney this range corresponds to a range of 199 $kgCO_2/m^2/year$ to 87 $kgCO_2/m^2/year$ after normalisation for hours effects. i.e. for otherwise equivalent buildings, there is a variation in actual emissions by more than a factor of two. Average performance in the market (which, being significantly dominated by VAV buildings, closely corresponds to the average performance of VAV buildings) is around 2.5 stars, or 151kg/m². This indicates that greenhouse abatements of up to 40% are available if the mid-range buildings can be made to perform as well as the best buildings.

With widespread tenant requirements for 4.5 star and higher rated buildings, the achievement of such savings is becoming a core business demand for building owners. Review of poorly performing buildings indicates that there are many common faults in VAV buildings, and that the majority of these can be tracked down to three basic parameters:

- **The underlying energy efficiency of the control routines.** This factor is particularly relevant in the Australian temperate climate, where buildings are designed for ambient temperatures typically in the region of 35°C but spend the majority of the year operating at 20-25°C. In this context, the operation of the building at low load is critical to energy efficiency.
- **The commissioning of systems to enable controls to function.** Commissioning practices within Australia are generally poor and, where energy efficiency is concerned, practically non-existent.
- **The poor robustness of the design in terms of energy efficiency.** Most VAV systems are built with multiple mechanisms for compensation, making the end result quite robust in terms of comfort. Unfortunately, these compensations generally cause additional energy consumption. As a result, standard VAV design is not robust in terms of energy efficiency. This is discussed in the following section in more detail.

The scale of these issues from an energy efficiency perspective is appreciable. Commercial experience indicates that for modern best practice VAV designs, achievement of 4-4.5 stars ABGR is generally practical once control and commissioning issues have been resolved. This suggests that the above factors may be related to greenhouse emissions penalties of 40% or more.

Design Considerations

The configuration of a standard VAV system is as follows. Air is heated or cooled in the central air handling plant. This air is distributed to variable air volume terminals, which consist of a damper which is used to regulate the quantity of air and, typically, a reheat element which is used to adjust the temperature of the air being delivered to the space. Air is then returned to the air handler, typically via return fan. Return air dampers are then used to adjust the degree to which air is recirculated. When external conditions are appropriate, full fresh air can be used providing what is variously known as free cooling or an air side economy or economizer cycle.

These elements combine to form a relatively flexible system. However, the relatively high component count leads to potentially complex control and maintenance issues. Prior to considering how these issues can be managed from a control perspective it is worthwhile to consider what can be done in terms of the basic VAV system design to make it more robust and reliable from an efficiency perspective[1].

Supply Air Volume

The total size of the system determines the upper limit on the fan energy that will be used. In the context that many VAV systems spend much of their life operating at well above 80% fan speed, typically due to poor commissioning, the maximum fan power is an important variable. The two key methods for managing this are:

- **Reducing fan size by lowering the supply temperature.** VAV systems are now being built with supply air temperatures as low as 7°C, thereby allowing fan sizes to be significantly reduced. This lowers the total fan power and thereby reduces the energy penalty if the fan control is poor, albeit with some penalty in chiller energy consumption.
- **Reducing the fan size through reducing downstream pressure losses.** The use of larger ducts and less flexible ducting can reduce fan power requirements.

A further point to note in this respect is that the use of large air handling equipment serving many floors or a very large area of building means that a single control or equipment fault can become an energy efficiency problem across a large part of the building. Thus there is an argument for reducing this risk by opting for smaller air handlers. The converse of this argument is that too many air handlers will create an excessive component count and thus a higher failure rate. These two considerations have to be balanced on a case-by-case basis.

Supply Air Temperature

The supply air temperature control is arguably the primary unifying variable of a VAV system. Therefore the effectiveness of this control depends on the level of load diversity served by the air handler. The more diversified the loads, the more likely that the peak load zone will require air that is excessively cool for other zones, causing excessive reheat operation. Thus by avoiding excessive thermal load diversity across the zones of an air-handler, the use of reheat can be minimised and the management of supply air temperature simplified.

---

[1]For a more thorough discussion of design issues see for instance [5] or [7].

A further common issue with supply air temperature control, from a hardware perspective, is the failure of valves in heating and cooling coils. In climates such as temperate coastal Australia, where the heating season lasts for only a few weeks per year, there is a good argument for the use of electric heating rather than hot water heating, purely because of the reliability issues and the avoidance of efficiency issues associated with low-load and sporadic operation inherent in a hot water system.

VAV Terminal Volume

Modern VAV terminals are designed to be capable of turning down to around 25% of the terminal unit's design supply volume. Below that level, the velocity sensing used to measure air flow generally becomes unreliable and poor control may result. From an energy efficiency perspective, the more a VAV unit can turn down, the better the outcome. However, a number of factors tend to cause turndowns in practice to be limited to 50% or poorer:

- **Oversizing of VAV terminals.** If VAV terminals are oversized, then the maximum required airflow is less than the terminal's capacity. As a result, once the new, restricted maximum is programmed, the terminal's minimum flow becomes a larger percentage of this revised maximum flow. This problem can be reduced by avoiding intentional oversizing and by rearranging zones to more closely match the available terminal unit sizes.
- **Air diffusion.** For standard four-way blow diffusers, even relatively limited reduction of airflow from design maximum can cause diffusion characteristics to become poor, leading to stratification in heating and "dumping" of cold air in cooling. This can be avoided through the use of higher performance diffusers such as swirl diffusers, light troffer air supply or linear diffusers.

VAV Terminal Reheat

VAV terminal reheat operation is a highly problematic issue for temperate climate zone VAV buildings. In a typical Australian VAV building, 10-15% of reheats will be operating because of some form of control, measurement or commissioning error, the most common of which tends to be the failure of the associated VAV terminal damper. This can constitute several hundred kW, which of course also creates a corresponding increase in chiller energy consumption. To avoid this, the following strategies can be used:

- **Avoidance of reheat.** If the system is well designed in terms of zoning and VAV terminal airflow turndown, it is generally possible to remove reheats from the system entirely. This solves the problem directly.
- **Limitation of reheat capacity.** If the system is designed to only provide trim heating at the terminal but has its primary heating capacity at the air handler, then the impact of reheat operation is greatly lessened due to the reduction in kW of reheat.
- **Withdrawal of reheats from terminals to branches.** If branch reheat is used, then this can be controlled from an aggregate of information from the terminals being served from that branch. This reduces the risk of operation to "cover" the failure of an individual terminal.

A further and important issue with respect to reheat is the use of hot water reheat. In a temperate climate such as Australia, the fundamental need for heating is low, so that much of the time the reheats should be off. This is easily achieved in an electric reheat system but can be highly problematic in a hot water system, as the valves may not always seal well. As a result, whenever the boiler system is running, there is a parasitic trickle of heat throughout the building from a number of reheats. This leads to a cascade of effects that dramatically reduce system efficiency. A single incorrectly operating reheat may drive the control zone temperature upwards, which will cause the supply air temperature to be reduced. This in turn may cause other zones to overcool and thus go into reheat operation. In this way a single component fault becomes a system-wide problem affecting all the zones served by the same air handling unit.

To avoid this, the simplest solution is not to use hot water reheat systems, and indeed it is notable that the buildings that this author has encountered with such systems have tended either to be poor performers or maintain efficiency by simply turning the boiler off for the vast majority of the year. This latter approach, while effective, calls into question the wisdom of the significant investment in hot water infrastructure in the original design.

Economizer

There are relatively few immediate design issues that arise with economizers, but the following factors are worth noting:

- The fresh air intake should be located so as not to be affected by building heat rejection systems;
- Separately controlled minimum fresh air and economizer cycle dampers should be used as it is difficult to obtain good control of minimum fresh air through the modulation of a full-air-flow damper set.

Control Approaches

In any real building, it is likely that there will be a combination of factors that prevent completely optimal design of the VAV air handling systems. Thus it becomes the role of the control system to operate in the presence of potential design weaknesses to reliably produce the most efficient outcome while still delivering comfort conditions. A range of approaches that are being used or advocated in Australia at present are discussed below.

Supply Fan Speed Control

Conventional practice speed control uses a variable speed drive to modulate the fan speed to maintain a fixed pressure set-point measured 2/3rds the way down the index run. In practice, this approach often leads to excessive fan operating speeds due to the following problems:

- **Poor commissioning of the supply air pressure set-point.** The commissioning engineer is typically tasked with finding a set-point that ensures that full flow can be provided to all terminals. The natural tendency in this situation is to select a set-point higher than the minimum required, to provide a “safety margin”.

- **Poor control philosophy.** The use of fixed static pressure control means that, at part load, the system operates with all VAV dampers closed to some extent thereby causing unnecessary throttling at terminal level. It has been identified that constant pressure systems tend to operate with an approximately quadratic variation in power with flow rate rather than with an exponent of 2.5 or higher for variable pressure systems [8].

In response to these challenges, it is necessary to control on a variable pressure basis. A number of methods for achieving this are listed below:

- **Critical zone reset.** In this approach, the fan speed is incrementally adjusted to keep one or more VAV dampers fully open. This is a relatively complex control, but is particularly useful with larger systems. It is important with this control to ensure that the selection of control zones takes account of the probability that 10-15% of terminals are in some form of failure mode. Furthermore, the conventional static pressure set-point should be retained as a maximum value to avoid over-pressuring caused by a failure in the control zone.
- **Crude pressure reset.** For small systems where there are relatively few VAV terminals, it is unnecessary to use individual VAV terminals as control zones. In this situation, either the fan speed itself can be directly adjusted in proportion to the total terminal demand flow, or the pressure can be reset in proportion to the total flow demand at the terminals. Both of these methods are crude and may starve individual terminals of airflow, but this is unimportant provided that the system serves largely similar zones that are conjoint.

Supply Air Temperature Control

The most common algorithms for supply air temperature control are based on a linear reset of the supply air temperature based on the temperature on a control zone. For instance, if the control zone is 24°C, then the supply air temperature is at minimum, while if the control zone is at 21°C the supply air is at maximum, with linear interpolation in between. The control zone in this instance is typically selected as the warmest zone served by the air-handler.

It may also be desirable for an additional reset to be applied to respond to reheat operation by increasing the set-point to ensure that reheat and fan energy are being appropriately balanced. The location of this balance will be dependent on a variety of building-specific factors.

For systems without reheat, a slightly different approach is required, as there is no means of controlling overcooling. In this case, the supply air reset is best operated on the basis of the average zone temperature. This will mean that it is possible that some zones overheat and some overcool, but in the absence of reheat, the best one can achieve is to approximately balance these two groups. This can be achieved through the use of a simple reset algorithm, possibly combined with a supplementary adjustment to actively balance the number of zones on either side of the control band.

The supply air temperature control is critical to determining the efficiency of the system in several ways, as evidenced by the following common failure modes:

- **Supply air temperature too low.** In this case the excessively low supply air temperature causes cooler zones to reach minimum comfort temperature and start overcooling, causing excessive reheat energy use.
- **Supply air temperature too high.** In this case the excessively high air temperature may cause the system to operate at high temperature and high volume, which is typically sub-optimal for energy efficiency.
- **"Neutral" temperature too low.** When the majority of the zones are in heating mode, it is not desirable for the primary air temperature set-point to be too low. A common control error in this respect is for the supply air temperature set-point when the control zone is at set-point to be too low. This then causes either the chiller(s) or the economizer to operate to reduce the supply air temperature, while in the majority of zones the reheat is operating. It is clearly preferable to be operating at a higher supply air temperature on minimum fresh air and without chiller operation under these conditions.

It can be seen that the relationship between supply air set-point, control zone conditions and terminal control activity is fundamental to system performance. To this end it is generally important to have a well defined relationship between the control zone temperature and the supply air temperature, such as achieved by the use of a reset. Thus, although more complex approaches can be used, such as a critical zone reset (which would, for instance, modulate the supply air temperature to the minimum possible while avoiding reheat use and retaining minimum airflow), the simple reset approach has significant merits.

It should also be clear from the above that the use of control zone temperature to directly control the cooling valve position is highly undesirable, as it does not provide any guaranteed relationship between zone temperature and supply air temperature. Furthermore, in the presence of over-sizing and variable air flow, this type of control often produces excessively low supply air temperatures, causing occupant comfort problems and excessive reheat operation.

A further issue to consider in this respect is the selection of a control zone. It is generally unwise to set the control zone to be the actual hottest zone for the system, as this zone is almost always in some form of failure mode. Preferred approaches are to select a control zone that is 10-15%, numerically, down the list of VAV terminals ranked in temperature order, or to base control zones on aggregated groups of VAV terminals.

VAV Terminal Control - Airflow

VAV terminals operate to increase flow in response to zone temperature. This is typically conceptualised as a proportional relationship where, as the zone temperature increases, the flow increases. In common with all forms of proportional control, however, this will mean that the temperature will increase in periods of high load. This can either be interpreted as a disadvantage, due to the potential loss of thermal comfort, or an advantage, in terms of providing some "float" in operating temperature that may mimic the adaptive response of occupants to warmer external conditions in general.

In practice, however, PI controls are often used to achieve a more exact set-point. The risks of such an approach are (a) the definition of excessively narrow control bands which may then cause increased heating/cooling conflict between adjacent zones and (b) the increase in overall energy use associated with the maintenance of lower loads in the presence of temperature dependent loads such as conduction through the building envelope.

These problems are illustrated in the following example from a single day of records from a single VAV terminal in an office building in Melbourne, as illustrated in Figure 1. It can be seen from the figure that the terminal has applied a wide range of airflows to keep the zone temperature with an exceptionally narrow (0.3°C) band. This has a number of potential negative consequences:

Figure 1. Airflow-Temperature Relationship for a PI Controlled VAV Terminal

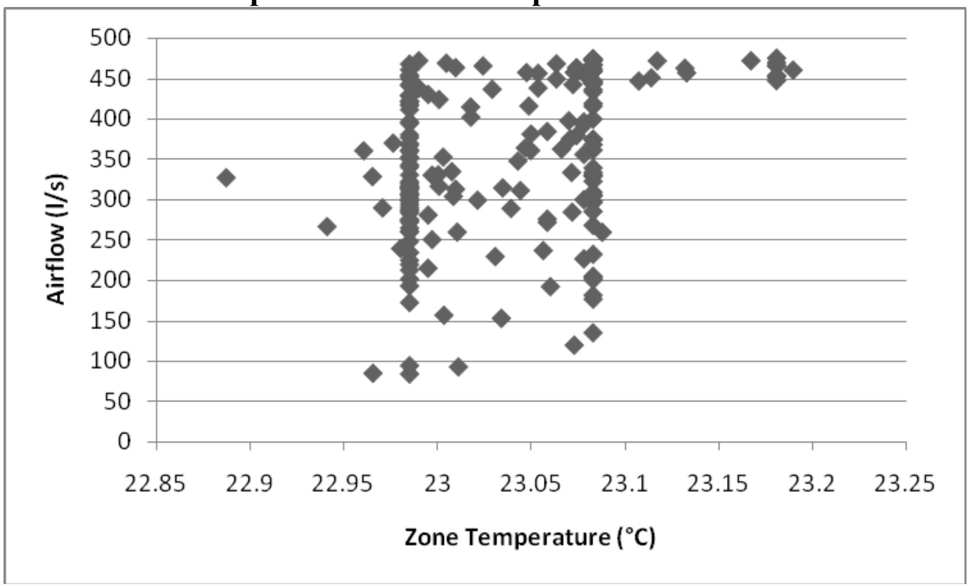

- The temperature is controlled within a narrow region, which will tend to increase the temperature dependent loads relative to a proportional control which would permit greater temperature drift.
- Given normal levels of variability in sensor calibration and representativeness, it creates a higher risk of undesirable interactions with adjacent zones.

The lack of a clear relationship between zone temperature and air flow makes the design of the design of the supply air temperature to zone temperature relationship more difficult. This is problematic as zone temperature is a natural and relatively stable proxy for zone load. By contrast, as can be seen in the It can be seen from the figure that the terminal has applied a wide range of airflows to keep the zone temperature with an exceptionally narrow (0.3°C) band. This has a number of potential negative consequences:

Figure 1, PI loop output can vary markedly over a very narrow temperature range, creating non-intuitive and potentially unstable control.

As a result, pure proportional control is strongly favoured. Successful configurations of this have typically been based on a 1°C deadband around the set-point, with the air volume increasing from minimum to maximum from set-point plus 0.5°C to set-point plus 1.5°C and the reheat operation operating between set-point minus 0.5°C to set-point minus 1.5°C. If reverse acting VAVs (i.e. VAVs where the air flow increases in heating mode as well as cooling mode) are used, then the volume increase below zone set-point should be disabled unless the supply air temperature is above the zone cooling set-point temperature. This is necessary because the volume increase below set-point would act to further cool the zone, and thereby increase reheat energy use (or, in the absence of reheat, further overcool the zone).

Other common errors for VAV terminal control include (a) too narrow a deadband between heating and cooling; (b) reverse action operation while the supply air temperature is cooling and (c) volume increase for cooling occurring through the deadband and/or the heating range.

VAV Terminal Control - Reheat

The management of reheat operation is of particular importance because of the role of reheat in "covering up" other failures in the system. Optimised control therefore should restrict reheat operation such that it only occurs when it is intended. In this context, simple thermostatic or proportional control at terminal level is inadequate as it does not reference the causal factors for the reheat demand. The two key options for management of reheat are:

- **Disabling the reheat of on an outdoor air lockout or similar.** This is a very effective way of preventing a great deal of parasitic reheat operation, as long as the system is designed to operate without reheat in cooling mode.
- **Enabling reheats on a group basis only.** Under this approach, groups of three to four related VAV terminals (for instance along a single facade of a single floor) are disabled from reheat until the average temperature of all the zones reaches conditions that merit reheat. This reduces the opportunity for an individual failed terminal to lapse into reheat.

In both cases, as with other measures that manage reheat, the consequence will be that it is more likely that a failed VAV terminal will generate an occupant complaint. While this is inconvenient for the building operator, it has the potential to alert them to the existence of a problem, with the result that it may be rectified.

Economizer Control

The principles of economizer control are relatively straightforward and as a result it is surprising how rarely it is well implemented. The basic principle is that the control should operate to enable the utilisation of outside air as a supplement to, or replacement of, the use of chillers.

For the outside air to be beneficial, it must have a lower energy content than the return air. This can be assessed on the basis of dry-bulb temperature and/or enthalpy. A well configured control has the following features:

- Economizer to be enabled when the outside air enthalpy is more than 2kJ/kg below the return air enthalpy (humid climates) and/or the outside air dry bulb temperature is below the return air dry bulb temperature (dry climates); and the outside air temperature is more than 1°C below the general zone set-point (this provides a failsafe in the case of return air sensor failure);
- The economizer should be disabled when the outside air enthalpy is equal to or above the return air enthalpy (humid climates) or the outside air dry bulb temperature is above either the return air dry bulb temperature (dry climates) or the general zone set-point.

The presence of the hysteresis in the above example is essential to prevent cycling between chiller and economizer operation.

A further opportunity as identified by [5,6] is in the sequencing of damper operation. Most buildings have the economizer programmed such that the return air damper closes and the outside air damper opens at the same rate as required to achieve the required supply air set-point. The same thermal result can be achieved if the outside air damper is fully opened prior to the return air damper closing, as this then creates a very low resistance airpath in the first stage of economizer operation.

Common errors on economizer programming are:

- **Use of relative humidity as a control criterion.** This is irrelevant and tends to reduce operation particularly in early morning operation when the relative humidity is high but the temperature and enthalpy are low enough to make economizer operation worthwhile;
- **Use of dry-bulb control in humid environments.** This can lead to significant loss of benefit through the increase in latent loads on the coils. Note that low temperature VAV systems (minimum supply air temperature below 10-12°C) also benefit from enthalpy or dewpoint control even in dry environments because of the need to prevent internal condensation.
- **Use of enthalpy control in dry environments.** As enthalpy sensors are a high failure item, if they are not delivering benefits then they should be avoided.
- **Use of single enthalpy sensors for control.** Calibration and failure issues with enthalpy sensors are common so it is preferable to work with some redundancy to allow for validation of enthalpy sensor readings. For buildings with large numbers of small VAV systems, it may be preferable to use a "deemed" return air enthalpy based on return air temperature and an estimated humidity (typically chosen low, for reasons of conservatism) to avoid the issues associated with the management of numerous humidity sensors.
- **Use of inappropriately low maximum dry-bulb limits.** It is common to find the maximum dry-bulb limit programmed to 18-20°C, thereby preventing operation across a wide range of potentially beneficial conditions.
- **Use of minimum dry-bulb limits.** It is reasonably common that minimum dry-bulb temperature limits are programmed, such that the economizer is disabled below typically 12-14°C outdoor air temperature. This is unnecessary as long as the supply air set-point control is working correctly. If this is a concern, a low-temperature lockout can be implemented that disables the economizer when the supply air temperature falls below the minimum supply air temperature.

Note that there are some variants worth considering for economizer operation. In climates with good economizer conditions, it is worth considering controlling the economizer to a different and more aggressive supply air temperature reset, so that when economizer is available it is used more aggressively than the chilled water valve is. This reduces chiller operation further than a conventional system, which would control both to the same temperature. In low temperature VAV systems, there is merit in considering a lockout of the chilled water valve and an upwards reset of the minimum supply air temperature to maximise economizer use when it can be demonstrated that this will still enable conditions to be met in the control zone. This can be assessed by comparing the energy flow in cooling at the current control zone flow and set-point with the energy flow available at the outside air temperature and full flow.

VAV System Control and Simulation

In the above, a wide range of control algorithms have been asserted as being good practice and certainly practical experience shows that such algorithms are capable of delivering building in the upper end of the efficiency spectrum. It is clear, however, that there are a number of significant judgements applied, including:

- The optimum balance between air volume and air temperature as determined by the detail of the supply air - control zone temperature relationship;
- The robustness or otherwise of the control to failures in control components;
- The merit of a range of modifications to the operation of the economizer.

In an ideal world, it would be desirable to test and optimise these via simulation. However, this is generally difficult because of the limitations of the control algorithms implemented within mainstream simulation packages, which are generally established to provide a working environment for plant but not intended as a means of detailed optimisation of this nature. However, in a market that is dominated by existing buildings, it is essential that the tools to optimise and advance control thinking in a fully integrated manner are available. In the presence of potentially large energy efficiency dividends, there is a strong case to be made for the improvement of control representation with the simulations to enable this type of optimisation to be undertaken.

For this to be achievable, the following key components need to be present in simulation packages' representation of VAV controls:

- Proportional control of VAV terminals;
- Ability to model systems with primary heating and no reheat;
- Switchable reverse action control based on supply air temperature;
- External lockout override of reheat operation;
- Reset of supply air temperature based on control zone temperature;
- Modifiable dynamic control zone selection (e.g. dynamically selecting the control zone but selecting the second warmest zone rather than the warmest);
- A variety of modifications to economizer operation including hysteresis, a separate economizer reset, and facility to optimise economizer operation against chiller operation.

A selection of these approaches is available in a number of packages, but no single package offers practical access to all of these options. This is a major issue as it means that the practical implementation of key algorithms is based on essentially little more than guesswork. The implementation of a realistic supply air temperature reset algorithm is the most critical item, as this standard control function cannot be properly represented or optimised in any of the major simulation packages to the knowledge of this author.

Furthermore, given the evident vulnerability of controls to component failures, consideration should also be given to the representation of such failures as a means of determining the robustness of the system. Key failures should include:

- Inaccuracy in zone temperature measurement
- Failure of VAV dampers
- Leaking of heat into the system via improperly sealed reheat valves;

Modifications of this type are recognised as being major new developments, but again, given the potential impact on the understanding of controls in the market, options of this kind merit consideration.

Conclusions

Variable air volume systems are a common and widespread air-conditioning system type, with a recognized potential to achieve good efficiency. However, work in Australia indicates that achievement of this potential is at best sporadic. Control and commissioning issues appear to dominate as factors in this disappointing outcome.

A range of design and control issues affecting VAV systems has been reviewed, looking at both the immediate energy efficiency impacts and the robustness to common component and commissioning failures.

The further optimisation of VAV controls requires deeper investigation than is currently possible via existing simulation programs. There is a strong case for future development of such programs to include upgrade of control representations to more closely match common industry practices to allow this optimisation to occur.

References

[1] For more information, see www.abgr.com.au

[2] For more information see www.energystar.gov

[3] "Australian Building Greenhouse Rating Scheme: Rating Operational Energy Efficiency in Office Buildings" P. Bannister and S. Salmon, Building Performance Congress, Frankfurt April 2004

[4] "Why Good Buildings Go Bad While Some Are Just Born That Way" P. Bannister, Australian Institute of Refrigeration, Air-Conditioning and Heating (AIRAH) Pre-Loved Building Conference, 28th November 2006, Melbourne Australia.

[5] Advanced Variable Air Volume System Design Guide, California Energy Commission, October 2003.

[6] ASHRAE Guideline 16-2003

[7] Variable Air Volume Systems for Environmental Quality, S. Chen, S Dempster, MacGraw-Hill, 1995

[8] Learning from Experiences with Energy Efficient HVAC Systems in Office Buildings, CADDET Analysis Series No. 15.