

This article was published in ASHRAE Journal, October 2011. Copyright 2011 American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. Posted at www.ashrae.org. This article may not be copied and/or distributed electronically or in paper form without permission of ASHRAE. For more information about ASHRAE Journal, visit www.ashrae.org.

High-Performance VAV Systems

By John Murphy, Member ASHRAE

Variable-air-volume (VAV) systems provide comfort in many different building types and climates. This article discusses several design and control strategies that can significantly reduce energy use in multiple-zone VAV systems. Although none of these strategies are new (in fact, several are required by energy standards or codes), implementation of them in buildings seems to be surprisingly infrequent. In this age of striving for higher levels of energy performance, these ought to be standard practice for every VAV system.

Optimized VAV System Controls

The first key ingredient to make a VAV system truly “high-performance” is the use of optimized system control strategies.¹

Optimal start/stop. Optimal start is a control strategy that uses a building automation system (BAS) to determine the length of time required to bring each zone from current temperature

to the occupied setpoint temperature. Then the system waits as long as possible before starting, so that the temperature in each zone reaches occupied setpoint just in time for occupancy (*Figure 1*).

This strategy reduces the number of system operating hours and saves energy by avoiding the need to maintain the indoor temperature at occupied set-

point even though the building is unoccupied.

Optimal stop is a control strategy that uses the BAS to determine how early heating and cooling can be shut off for each zone so that the indoor temperature drifts only a few degrees from occupied setpoint before the end of scheduled occupancy (*Figure 1*). In this case, only cooling and heating are shut off; the supply fan continues to operate and the outdoor-air damper remains open to continue ventilating the building.

This strategy also reduces the number of system operating hours, saving energy by allowing indoor temperatures to drift early.

Fan-pressure optimization. As cooling loads change, the VAV terminals modulate to vary airflow supplied to the zones. This causes the pressure inside the supply ductwork to change. In many systems, a pressure sensor is located

About the Author

John Murphy is an applications engineer with Trane, a business of Ingersoll Rand, in La Crosse, Wis.

approximately two-thirds of the distance down the main supply duct. The VAV air-handling (or rooftop) unit varies the speed of the supply fan to maintain the static pressure in this location at a constant setpoint. With this approach, however, the system usually generates more static pressure than necessary.

When communicating controllers are used on the VAV terminals, it is possible to optimize this static pressure control function to minimize duct pressure and save fan energy. Each VAV controller knows the current position of its airflow-modulation damper. The BAS continually polls these individual controllers, looking for the VAV terminal with the furthest-open damper (Figure 2). The setpoint for the supply fan is then reset to provide just enough pressure so that at least one damper is nearly wide open. This results in the supply fan generating only enough static pressure to push the required quantity of air through this “critical” (furthest-open) VAV terminal.

At part-load conditions, the supply fan is able to operate at a lower static pressure, consuming less energy and generating less noise.

Fan-pressure optimization provides the added benefit of allowing the building operator to identify and address “rogue zones.” A rogue zone is one where something is not working properly. Some possible causes include an undersized VAV terminal, a restriction in the duct that does not allow the required airflow, a zone temperature setpoint that has been reset too low, or a zone sensor installed in the sunlight or near a heat source, like a coffee maker.

Whatever the cause, with conventional constant duct pressure control, the building operator only learns about these problems when someone complains about comfort or noise. However, with fan-pressure optimization, the BAS regularly gathers data from each VAV terminal, providing the opportunity to identify and fix these rogue zones.

Figure 3 includes a chart that trends the position of VAV dampers in an actual building. In the morning, most of the dampers open a little after 7 a.m. to warm up the zones, but then they close down quite a bit, indicating that not much cooling is needed on this day. The VAV terminal serving Room 204, however, is nearly wide open for most of the day, and is preventing the duct pressure setpoint from being reset downward to reduce fan energy.

This suggests that there may be a problem with this zone. During the first few months of operation, the building operator can review this trend periodically, identify any potential rogue zones, and add them to his “to be fixed” list. If the BAS has the capability to temporarily exclude a rogue zone from the control sequence, it would allow that VAV




Figure 1: Optimal start/stop.




Figure 2: Fan-pressure optimization.

terminal to operate as normal and attempt to control zone temperature, but would not have a “vote” when determining the optimized duct static pressure setpoint. After a while, all rogue zones would be identified and the fan-pressure optimization sequence should be operating without one dominant, or rogue, zone. At that time, a technician can be dispatched to fix all the rogue zones and re-include them in the sequence.

Supply-air-temperature reset. In a VAV system, it is tempting to raise the supply-air temperature (SAT) at part-load conditions to save compressor and/or reheat energy. Increasing the SAT reduces compressor energy because it allows the compressors to unload or cycle off. In addition, SAT reset makes an airside economizer more beneficial. When the outdoor air is cooler than the SAT setpoint, the compressors are shut off, and the outdoor- and return-air dampers modulate to deliver the desired supply-air temperature. A warmer SAT setpoint allows the compressors to be shut off

sooner and increases the number of hours when the economizer is able to provide all the necessary cooling.

For zones with low cooling loads, when the supply airflow has been reduced to the minimum setting of the VAV terminal, raising the supply-air temperature also decreases the use of reheat at the zone level. However, because the supply air is warmer, those zones that require cooling will need more air to satisfy the cooling load. This increases supply fan energy.

Finally, in climates that experience humid weather, warmer supply air means less dehumidification at the coil and higher humidity levels in the zones. If dehumidification is a concern, use caution when implementing this strategy.

SAT reset should be implemented so that it minimizes overall system energy use. This requires considering the trade-off between compressor, reheat, and fan energy, while not ignoring space humidity levels. Although there are several different approaches to implementing this strategy, many of the papers and articles written on the subject tend to agree on some general principles for balancing these competing issues.^{1,2,3}

First, when it is warm outside, keep the SAT cold. This takes advantage of the significant energy savings from unloading the fan. Then, begin to raise the SAT setpoint during mild weather when it can enhance the benefit of the airside economizer and reduce reheat energy.

Although there are several possible control sequences, the example in *Figure 4* demonstrates one way to attempt to balance these competing impacts of SAT reset.

With this approach, the SAT setpoint is reset based on the changing outdoor dry-bulb temperature. When the outdoor dry-bulb temperature is warm—higher than 65°F (18°C) in this example—no reset takes place and the SAT setpoint remains at its design value of 55°F (13°C). When it is warm outside, the outdoor air provides little or no cooling benefit for economizing, and the cooling load in most zones is likely high enough that reheat is not required to prevent over-cooling. Keeping the air cold allows the fan to turn down, taking advantage of the energy savings from reducing airflow. In addition, the colder supply-air temperature allows the system to provide sufficiently dry air to the zones, improving part-load dehumidification.

When the outdoor temperature is cooler, the controls begin to reset the SAT setpoint upward. At mild or cold outdoor temperatures, reset enhances the benefit of the economizer, and if there is any zone-level reheat, it is reduced or even avoided. At these cooler temperatures, the supply fan has likely already unloaded significantly, so the incremental energy use of having to deliver a little more air is lessened.




Figure 3: Identifying rogue zones.




Figure 4: Example of supply-air temperature (SAT) reset control.

Finally, the amount of reset is limited, to 60°F (16°C) in this example, which allows the system to satisfy cooling loads in interior zones without needing to substantially oversize VAV terminals and ductwork.

A possible drawback of this approach is that if a zone has a high cooling load, even when it is cool outside, it is possible that the warmer supply air may not provide enough cooling and that zone may overheat. To prevent this from occurring, when SAT reset does take place, the amount of reset should depend on the cooling need of the worst-case zone. (Zones that are expected to have near-constant cooling loads should be designed for the maximum reset SAT.)

As depicted in the example in *Figure 4*, when it is 50°F (10°C) outside, the system will attempt to raise the SAT setpoint to 60°F (16°C). However, if there is a zone that is at near-design cooling load, this air may not be cold enough. Using the current temperature and VAV damper position for that

Advertisement formerly in this space.

Optimal start (Section 6.4.3.3.3) and **zone-level demand-controlled ventilation** (Section 6.4.3.9) are mandatory requirements of ASHRAE/IES Standard 90.1-2010.⁶ **Fan-pressure optimization** (Section 6.5.3.2.3), **supply-air-temperature reset** (Section 6.5.3.4), and **system-level ventilation reset** (Section 6.5.3.3) are prescriptive requirements of the standard.

zone, the controls can determine that this amount of reset is too much, and either change the SAT setpoint back to 55°F (13°C), or lower the SAT setpoint a degree or so and see if that eliminates the overheating problem.

Ventilation optimization. In a typical VAV system, the VAV air-handling (or rooftop) unit delivers fresh outdoor air to several, individually controlled zones. Demand-controlled ventilation (DCV) involves resetting intake airflow in response to variations in zone population. Although commonly implemented using carbon dioxide (CO₂) sensors, occupancy sensors or time-of-day (TOD) schedules can also be used.

One approach to optimizing ventilation in a multiple-zone VAV system is to combine these various DCV strategies at the zone level (using each where it best fits) with ventilation reset at the system level.

In the example system depicted in *Figure 5*, CO₂ sensors are installed only in those zones that are densely occupied and experience widely varying patterns of occupancy (such as conference rooms, auditoriums, or a lounge area). The VAV controller resets the ventilation requirement for that zone based on the measured CO₂ concentration.

Zones that are less densely occupied or have a population that varies only a little (such as private offices, open plan office spaces, or many classrooms) are probably better suited for occupancy sensors. When a zone is unoccupied, the VAV controller lowers the ventilation requirement for that zone, typically to the building-related ventilation rate, R_a , required by ASHRAE Standard 62.1.

Finally, zones that are sparsely occupied (such as open office areas) or have predictable occupancy patterns (such as cafeterias or auditoriums) may be best controlled using a time-of-day schedule. This schedule can either indicate when the zone will normally be occupied versus unoccupied, or can be used to vary the ventilation requirement based on anticipated population of that zone.

These various zone-level DCV strategies can be used to reset the ventilation requirement for their respective zones. This zone-level control is then tied together using ventilation reset at the system level (*Figure 5*). In addition to resetting the zone ventilation requirement, the controller on each VAV terminal continuously monitors primary airflow being delivered to the zone.

The BAS periodically gathers the current ventilation requirement and primary airflow from all the VAV control-




Figure 5: Ventilation optimization (DCV at zone level + ventilation reset at system level).

lers and solves the ventilation reset equations (prescribed by ASHRAE Standard 62.1⁵) to determine how much outdoor air must be brought in through the system-level OA intake to satisfy all zones served. Finally, the BAS sends this outdoor airflow setpoint to the VAV air-handling (or rooftop) unit, which modulates a flow-measuring outdoor-air damper to maintain this new ventilation setpoint.

“Occupied standby” mode. When an occupancy sensor is used in combination with a time-of-day schedule, the sensor can be used to indicate if the zone is unoccupied although the BAS has scheduled it as occupied. This combination can be used to switch the zone to an “occupied standby” mode.

In this mode, all or some of the lights in that zone can be shut off, the temperature setpoints can be raised or lowered by 1°F to 2°F (0.5°C to 1°C), and the ventilation requirement for that zone can be reduced, typically to the building-related ventilation rate, R_a , required by Standard 62.1.

In addition, for a VAV system the minimum airflow setting of the VAV terminal can be lowered to avoid or reduce the need for reheat. This minimum airflow setting is typically set to ensure proper ventilation. However, when nobody is in the room and with the ventilation requirement reduced, the minimum airflow setting can be lowered significantly during this occupied standby mode. This reduces both reheat and fan energy use.

When the occupancy sensor indicates that the zone is again occupied, these settings are switched back to normal occupied mode.

Variable airflow during heating. The conventional way to control a VAV reheat terminal has been to reduce primary airflow as the zone cooling load decreases. When primary airflow reaches the minimum setting, and the cooling load continues to decrease, the reheat coil is activated to warm the air and avoid overcooling the zone. When this occurs, the airflow-modulation damper typically maintains a constant heating airflow.

Advertisement formerly in this space.

Figure 6 depicts an alternate method to control a VAV reheat terminal. When the zone requires cooling, the control sequence is unchanged; primary airflow is varied between maximum and minimum cooling airflow as needed to maintain the desired temperature in the zone.

When primary airflow reaches the minimum cooling airflow setting, and the zone temperature drops below the heating setpoint, the heating coil is activated to warm the air to avoid overcooling the zone. As more heat is needed, the controller resets the discharge-air temperature setpoint upward to maintain zone temperature at setpoint (orange dashed line in Figure 6), until it reaches a defined maximum limit—90°F (32°C) in this example. The discharge temperature is limited to minimize temperature stratification when delivering warm air through overhead diffusers.

When the discharge-air temperature reaches this maximum limit and the zone requires more heating, primary airflow is increased while the discharge-air temperature setpoint remains at this maximum limit. The result is that the airflow-modulation damper and hot-water valve will modulate open simultaneously.

By actively controlling the discharge-air temperature, it can be limited so that temperature stratification and short circuiting of warm air from supply to return are minimized when the zone requires heating. This improves occupant comfort and results in improved zone air-distribution effectiveness, which avoids wasteful over-ventilation.

Note: Section 6.5.2.1 of Standard 90.1-2010 allows this alternate control strategy (as long as the maximum heating primary airflow is less than 50% of maximum cooling primary airflow) as an exception to comply with the Standard's prescriptive limitation on simultaneous heating and cooling.

Cold-Air Distribution

Another key ingredient of some high-performance VAV systems, especially chilled-water VAV systems, is lowering the supply-air temperature.^{1,4}

Supplying air at a colder temperature—between 45°F to 52°F (7°C to 11°C) for example—allows the system to deliver a lower supply airflow rate. This can significantly reduce fan energy use, and it can also allow fans, air-handling units, and VAV terminals to be downsized, which reduces installed cost. Sometimes, ductwork is downsized also, which further reduces installed cost.

Another potential benefit is that delivering colder air means that the air is drier, which can lower indoor humidity levels in climates that experience humid weather.

Although supplying air at a colder temperature reduces fan energy use, it requires the use of a colder chilled-water temperature (which impacts chiller efficiency), increases




Figure 6: Control of a VAV reheat terminal to vary airflow during heating.

reheat energy, and results in fewer hours when the airside economizer can provide all the necessary cooling. Although the lower humidity levels that result from colder air may be appreciated in some applications, this “extra” dehumidification results in an increased latent load on the chiller. This increased latent load is partially offset by a reduction in the sensible load due to fan heat (reduced fan power). These impacts partially offset the fan energy savings. Therefore, whole-building energy simulation should be used to determine the impact of a lower supply-air temperature on the overall energy use of a VAV system.

The first tip to maximizing energy savings in a cold-air VAV system is to reset the SAT setpoint upward during mild weather. As explained previously, this helps maximize the benefit of the airside economizer and reduces reheat energy use, while still achieving fan energy savings during warm weather.

The second tip is to try raising the zone temperature setpoint by one or two degrees. Since people are comfortable at warmer temperatures when humidity is lower, the lower humidity levels that occur with cold-air systems provide the opportunity to slightly raise the zone cooling setpoint. This further reduces airflow and fan energy use, and reduces cooling energy a little too.

Next, while lowering the supply-air temperature can allow the ducts to be downsized to reduce installed cost, if a goal of the project is to maximize energy savings, consider designing the system for the colder supply-air temperature but not downsizing the ductwork as much as possible. Keeping the ducts a little larger reduces fan energy and allows SAT reset to be used without concern for any zones with near-constant cooling loads. It also improves the ability of the air distribution system to respond to any future increases in load, since it will be capable of handling an increased airflow rate if needed. Use caution when oversizing VAV terminals to ensure that they are able to properly operate across the entire range of expected airflows.

Challenges. Of course, cold-air VAV systems are not without challenges. Design engineers typically express two con-

Advertisement formerly in this space.

cerns with this approach: 1) minimizing comfort issues related to cold air dumping on the occupants, and 2) avoiding condensation on components of the air distribution system. A great resource for anyone designing a cold-air system is the *ASHRAE Cold Air Distribution System Design Guide*, which discusses in detail how to avoid these problems.⁴

Linear slot diffusers with a high induction ratio are good for any VAV system, but work very well for a cold-air VAV system. They induce large quantities of air from the space to mix with the cold supply air, so the mixture drops slowly into the occupied zone at nearly room temperature. On the other hand, many conventional diffusers allow cold air to drop directly into the occupied zone, which can result in occupant comfort complaints, especially at reduced airflows.

Since the surfaces of the air distribution components in a low-temperature system are colder than in a conventional system, there is often concern about condensation. To minimize the risk of condensation:

Advertisement formerly in this space.




Figure 7: Example energy analysis for a large office building.

- Insulate all the cold surfaces, including supply ducts, VAV terminals, and diffusers. In addition, a properly sealed vapor retarder should be included on the warm side of the insulation to prevent condensation within the insulation.

- If possible, use an open ceiling plenum for return air. This results in a conditioned plenum, which means a lower dew point and less risk of condensation.

- During humid weather, maintain positive building pressurization to reduce or eliminate the infiltration of humid outdoor air.

- During startup, slowly ramp the SAT setpoint downward to slowly lower surface temperatures while lowering the dew point inside the building.

Putting It All Together

Although this article focused on the airside of VAV systems, many opportunities exist to reduce energy on the equipment or plant side of the system. Some examples include:

- High-efficiency rooftop equipment, water chillers, and air-handling unit fans;
- Air-to-air energy recovery;
- Evaporative condensing;
- Low-flow, low-temperature chilled-water systems;
- Waterside heat recovery;
- Central geothermal; and
- Solar heat recovery for reheat.

The impact of any of these strategies on overall operating costs depends on climate, building use, and utility costs. Therefore, whole-building energy simulation should be used to determine if a specific strategy makes sense for a given application. As an example, *Figure 7* contains the results from a whole-building energy simulation of a large office building, comparing a typical VAV system to a high-performance chilled-water VAV system.

The baseline building uses a conventional chilled-water VAV system, designed for 55°F (13°C) supply air, and modeled according to Appendix G of Standard 90.1-2007. The high-performance VAV system is designed for 48°F (9°C) supply air (ductwork has not been downsized) and uses the op-

Advertisement formerly in this space.

timized VAV control strategies mentioned in this article, and a low-temperature, low-flow water-cooled chiller plant.

For this example, the building with the high-performance VAV system uses about 20% less energy than the baseline VAV system in Houston, Philadelphia, or St. Louis. The building uses about 10% less in Los Angeles, which has milder weather and lots of hours for airside economizing. As a comparison, the high-performance VAV system uses 5% to 10%

less energy than an active chilled beam system that was modeled for this same building.

Finally, over the past several years, a series of *Advanced Energy Design Guides*, have been jointly developed by the U.S. Department of Energy, ASHRAE, the American Institute of Architects, the Illuminating Engineering Society, and the U.S. Green Building Council.⁷ These guides include climate-

specific recommendations that can be used to achieve 30% (or in some cases 50%) energy savings over conventional design. Seven guides are currently available in this series, covering building types from office buildings to schools to warehouses.

Most of these guides include several options for HVAC systems. In several of them, VAV systems are one of the options covered that can help the overall building achieve the stated energy-savings threshold. For example, in the recently published guide for small- and medium-sized office buildings, a high-performance rooftop VAV system is included as one of the options that can be used to achieve 50% energy savings. In the guides for K-12 school buildings and small health-care facilities; rooftop VAV and chilled-water VAV systems are included as options for achieving 30% energy savings.

These guides, and the whole-building energy simulations that were used to confirm the climate-specific recommendations contained in them, provide validation that VAV systems can be used in high-performance buildings.

Advertisement formerly in this space.

References

1. Murphy, J., B. Bakkum. 2009. *Chilled-Water VAV Systems*. La Crosse, Wis.: Trane.
2. Energy Design Resources. 2009. *Advanced Variable Air Volume System Design Guide*. Sonoma, Calif.
3. Wei, G., M. Liu, D. Claridge. 2000. "Optimize the supply air temperature reset schedule for a single-duct VAV system," *Proceedings of the Twelfth Symposium on Improving Building Systems in Hot and Humid Climates*.
4. ASHRAE. 1996. *Cold Air Distribution System Design Guide*.
5. ANSI/ASHRAE Standard 62.1-2010, *Ventilation for Acceptable Indoor Air Quality*.
6. ANSI/ASHRAE/IESNA Standard 90.1-2010, *Energy Standard for Buildings Except Low-Rise Residential Buildings*.
7. ASHRAE. *Advanced Energy Design Guide* series. www.ashrae.org/technology/page/938.●