MEASURED ENERGY CONSUMPTION OF VARIABLE-AIR-VOLUME FANS UNDER INLET VANE AND VARIABLE-SPEED-DRIVE CONTROL

D.M. Lorenzetti

L.K. Norford, Ph.D.
Member ASHRAE

ABSTRACT

Variable-speed drives (VSDs) are a popular retrofit for variable-air-volume ventilation fans, yet few measurements of annual energy savings have been performed. To eliminate the need for a year of pre-retrofit and a year of post-retrofit data, one previously developed measurement technique (applied to two air handlers in a single building) relied on correlations of fan power to volumetric flow rates combined with histograms of flow data. There is a need for more performance measurements, particularly when flow data are not available.

To the extent that the demand for airflow depends on thermal gains and building occupancy and use patterns, fan power may be correlated to these directly. One advantage to this approach is that knowledge of temperature and solar loading profiles, as well as expected occupancies, may be obtained more easily than is the case for airflow. In this paper, fan power is modeled as a function of outside temperature only for four variable-air-volume systems outfitted with VSD controllers. The curves developed are used to estimate the savings of VSDs over variable inlet vanes in the same air handlers. Practical problems, such as filtering data and the choice of sampling intervals, are discussed.

INTRODUCTION

Variable-speed motor drives (VSDs) are an increasingly popular technique for energy-efficient ventilation of commercial buildings. VSDs are significantly more expensive than the variable inlet vanes (VIV) they replace, and many building owners and operators turn to electric utility conservation assistance programs to fully or partially fund VSD retrofits or installation in new construction. While financial arrangements vary, a utility might offer a customer a rebate on the basis of projected savings, and the customer enters into a shared or guaranteed savings contract with an energy service company, which performs the work and is compensated from expected savings. All parties therefore have a stake in the savings, which typically are estimated on the basis of manufacturers' literature. Because of the expense and prevalence of VSD technology, there is an obvious need for measured savings data, both to validate and to improve the estimations. Until computational procedures are shown to be sufficiently accurate, or measurement data are available from a sample of buildings large enough to make inferences about the population as a whole, there is a need for simple yet reliable techniques to make these measurements.

Both data and measurement methodologies for VSD fan motor retrofits have been investigated by Englander and Norford (1992a, 1992b). The data came from two supply fan/return fan air-handling units in a commercial building. Pre-retrofit data, under VIV control, were taken for a one-year period. Hourly average electricity use for the four fans was correlated with measured airflows. Measured annual electricity use was closely approximated by a convolution of the power versus flow correlation and a histogram of hourly average airflows. After the retrofit, the same type of correlation was established with VSDs over a period of about six weeks and applied to the base-year flow histogram to project annual post-retrofit consumption. Savings were 12% and 45% for the two supply fans, 56% and 67% for the two return fans, and 35% in aggregate. When the static pressure setpoint for the two supply fans was lowered from 2.5 in. w.g. to 1.5 in. w.g. for a period of about two weeks, the projected annual savings increased to 46% and 67% for the supply fans and 56% in aggregate. Relatively lower savings for one of the supply fans was attributed to high airflows demanded by poorly maintained VAV terminal boxes in an unoccupied part of the building.

The methodology reported by Englander and Norford offered the advantage of isolating the measurements to the fan itself and relying on a correlation that, while empirical, can be traced to the physics of air movement and the efficiency of the fan and motor. The power-airflow correlation is used in many energy analysis codes, including DOE-2 and bin methods. The measured correlation appeared to be relatively free of unexplained influences on power, with some slight data scatter due to deviations in static pressure from the setpoint.

David M. Lorenzetti is a research assistant and **Leslie K. Norford, Ph.D.**, is an assistant professor of building technology, Department of Architecture, Massachusetts Institute of Technology, Cambridge.

However, the strength of the correlation depends on a somewhat difficult and expensive measurement of airflow. Some air-handling units use airflow sensors, typically single- or multiple-point Pitot tubes, to control the return fan in response to changes in supply airflow. Other systems, such as the four presented here, require no flowmeters and rely instead on open-loop synchronization to control the return fan as well as the supply fan from the same signal, derived as the difference between measured and desired static pressure. Installation and maintenance costs influence decisions about the use of airflow sensors, but even with such sensors installed as part of the HVAC system, and thus available at no cost for energy conservation measurements, the cost of acquiring a year of flow data remains.

This paper presents energy consumption data for four supply fans installed in a university building that houses a mix of teaching, office, and clinic spaces associated with a health sciences program. The methodology, based on correlating fan power with outside temperature, develops a weaker relationship than with airflow, but one that can be combined with readily available local temperature data to significantly reduce the cost of information.

DESCRIPTION OF VENTILATION AND MONITORING SYSTEMS

Six air handlers and two additional lab-hood makeup air fans (a total of 14 fans) in two connecting buildings were retrofitted two years ago with pulse-width-modulated VSDs. Four of these air handlers are in a single mechanical room, where the data collection efforts are concentrated. The inlet vanes for all fans were left intact during the retrofit and normally remain fully open. The control system for each fan includes switching that permits the VSDs to be shut down and control to revert to inlet vanes, a feature that facilitates VSD maintenance. This feature allowed us to compare inlet vane and VSD performance in a "flip flop" mode, in lieu of the usual "before and after" sequence that was not possible given the timing of our involvement in the project. We do not regard this as a detriment, for switching control modes allows focused data collection under the operating conditions (outside temperatures) of interest, a strategy that is more time-efficient than the sequential approach. This strategy also improves the statistical quality of the data by blocking against the tendency of events such as filter changes, load rebalancing, and corrective maintenance to alter the system's long-term characteristics (Box et al. 1978).

The motors powering the fans under study range from 7.5 to 40 hp (see Table 1). The last column of Table 1 shows the fan power needed to provide the rated airflow at design pressures, as determined from the manufacturers' fan curves. The motor for supply fan 9 in particular appears to be substantially oversized.

Fan power was monitored with a data logger designed to reduce the cost and complexity of recording electrical power. For three-phase/three-wire motors, two current transducers are required. Currents are sampled at 0.25-second intervals and multiplexed with a voltage tap that reads potential from each phase. The logger computes true power and stores hourly average readings in memory for subsequent downloading. In addition to power data, each hourly record includes the average outside temperature and the temperatures in two of the supply ducts, measured using resistive temperature devices. Airflow data, not used for this study, were taken for two fans by reading the electrical output of pressure transducers connected to Pitot arrays no longer used by the fan controllers. In this paper, we analyze data on the four supply fans taken from August 1990 to April 1991.

CORRELATING FAN POWER WITH OUTSIDE TEMPERATURE

Models

A variable-air-volume (VAV) ventilation system controls airflow in response to cooling load. Buildings

TABLE 1
System Characteristics for the Four Air Handlers

Fan	Floor area, 1000 ft^2	cfm x1000	cfm/ ft^2	Motor horse-power	hp/ kcfm	bhp at rated cfm
AH-8-SF	14.9	21.5	1.4	25	1.2	21
AH-8-RF	14.9	16.0	1.1	7.5	0.5	
AH-9-SF	20.1	27.2	1.4	40	1.5	26
AH-9-RF	20.1	20.6	1.0	10	0.5	
AH-10-SF	16.9	15.0	0.9	40	2.7	28
AH-10-RF	16.9	12.8	0.8	7.5	0.6	
AH-11-SF	30.9	32.9	1.1	40	1.2	36
AH-11-RF	30.9	21.5	0.7	15	0.7	

typically show measurable changes in cooling load as the outside temperature varies, although these changes may be small compared with constant internal loads and may be masked by fluctuating internal and solar gains, by changes in thermostat setpoints, and by deviations in zonal temperatures around the setpoint. It is nonetheless reasonable to expect some degree of correlation of fan power with outside temperature. Should the ventilation system operate under systematic changes in internal loads and thermostat values (for example, under a night setback scheme), separate temperature correlations can be performed over each period.

For these fans, plots of power against outside temperature suggest a piecewise fit in which power is constant below some breakpoint temperature and is increasing above. This breakpoint marks the power needed to establish some minimum airflow determined by the flow requirements in individual offices. Below it, further reductions in flow are prohibited.

While higher-order piecewise fits are possible, in this paper we consider piecewise linear fits only. Partly this is because the piecewise linear fit has been the model of choice for studies of whole-building energy consumption, particularly of residential energy use, and is often based on monthly billing data (Fels 1986). In large part, however, it also is due to the fact that adding terms beyond the linear did not in many cases allow the model to explain any more of the variance in the data. Statistically, adding parameters shifts degrees of freedom from the data to the model. When the greater latitude in the model is not matched by a greater power to explain the variation in the data, the added parameters have weakened the model. Mathematically, this occurs when the additional parameters do not decrease the error sum of squares enough to compensate for the smaller variance attributable to the model, which varies roughly with the inverse of the number of parameters.

This same effect was noted for models that attempted to incorporate some knowledge of the order in which the power-temperature pair was recorded during the day. Adding a model parameter to account for the relative time in the day did reduce the sum of squares associated with the error terms but not by enough to offset the fact that, all else being equal, added model parameters reduce the variance explained by a model. Thus, the models considered in this paper all are two- and three-parameter fits to the data. They are the piecewise linear fit, a simple linear regression, and two quadratic models (one with and one without a linear term).

Each model, by taking power to be a function of outdoor temperature without regard for the dynamics of the building's control system, thermal mass, internal gains, and so on, implicitly represents a steady-state system. This assumption breaks down at time intervals that are short relative to the thermal and controller time constants. Fortunately, hourly average data rarely are affected by controller dynamics (an exception is during prolonged changes in airflow, when fan power is influenced by controller proportional gain and attendant errors in the pressures against which the fan operates). Building thermal time constants, on the other hand, are much longer, introducing scatter to simple correlations with outside temperature—especially in cases where an air handler is shut down overnight, allowing the building temperature to "float" until a thermal pulldown cycle begins the next morning.

Dynamic models can account for thermal mass but, by requiring a time history of temperatures, they remove the potential simplification of combining correlated powers with yearly temperature histograms. Alternatively, one may suppress the dynamics by averaging power and temperature data over longer intervals, e.g., daily. Therefore, in addition to hourly data, the four models were fit using daily averages and daily averages weighted for hours of operation.

Conditioning the Data

Given a data file containing a time series of hourly averages of power and temperature, a typical set of preprocessing steps involves separating the records between VSD and VIV modes, filtering for inconsistent data, and withholding a number of records for use as a quasi-independent check on the results. Finally, the daily averages are calculated and the respective data files are converted into a format compatible with the statistical analysis package.

Of these tasks, the most challenging numerically is the filtering step. Specifically, only one of the four air handlers runs 24 hours a day, while the rest run for anywhere between 8 and 15 hours a day. Therefore, two hourly records a day are averages of both on- and off-power readings. These transition points, because they violate the steady-state assumption, must be filtered out of the data set before it is used to derive the models. In addition, the power drawn just after start-up may be abnormally high due to the building pulldown mentioned above. If so, some thought should be given to discarding these points as well. Because our aim was to predict total fan power, we did not remove start-up surges from the data sets. The filtered points, however, cannot be discarded entirely, since they must be used later to verify that the models estimate total fan energy accurately.

Three filtering strategies were considered. All depend on knowing the time order of the power records, and two require that some nominal power be established below which the fan may be considered shut off. The third uses a student's T-distribution to detect changes in fan power and is only made more robust by the introduction of cutoff power.

The most naive filter presupposes an abundance of data; it simply strikes the first and last entries from any sequence of power-on records. When checking total fan energy, we counted half the energy the model gave for filtered points on the assumption that on average, they would combine half an hour of running and half an hour of shutdown information. Therefore, using this first method can cause underestimates of total fan energy, since more points are filtered than might be the case using less arbitrary criteria. The advantage of this method is that the only way to miss a mixed-power point is if the fan shuts down for so short a time that the defined fan-off power level never is reached. This actually happens, on occasion, for fan AH-11-SF, generally between 7:00 p.m. and 8:00 p.m., when the power falls off dramatically but not enough to trigger the filter. The only guard against this sort of behavior is to graph the data and investigate the nature of disturbances in time and with temperature.

The second filter tests each transition point against its nearest power-on neighbor, comparing their ratio to some critical fractional power. Less heavy-handed than the first, this fractional test is simple to implement and acts reliably. Finally, the student's T-test is the most sophisticated, and generally the most predictable, of the filtering rules but breaks down if the fan on-off states are not clearly maintained for longer than three hours, the minimum length of time needed to establish a reference distribution of powers for the T-test.

The complexities and uncertainties of filtering could have been avoided by the provision of a logic input to the data logger from the fan controller itself. This input would be high when the fan was energized and low when the fan was off. Then each power record would include, in addition to the average power and average temperature read over the course of the hour, an entry showing the fraction of that hour during which the fan had run. Given a logger able to average digital inputs over time, a natural weighting scheme would arise by which to give each datum more or less force in determining the power-temperature relationship. The value of knowing how long the fan was on and off during these mixed-power hours should not be underestimated when choosing and configuring a logger.

Once filtered to remove transition points between on- and off-states, a number of data records were stripped away for later use in checking the models' predictive ability. To achieve a uniform distribution of sampled points, it is important to avoid selecting them from groups made up of factors of 24 (hours in a day) and 22 (hours in a day lacking two filtered points) or multiples of 7 (days in a week). For example, taking every sixth record from the file risks sampling only four distinct hours over long runs of time. This risk is less clear once a file has been filtered to remove power-off and mixed-power points, but in general one must be wary of choosing sampling periods that coincide with the natural period of the file. These considerations in mind, we took 5, 13, 17, and 19 to be appropriate sampling periods. Finally, to assemble a uniformly distributed sample within less than half a month of data, we would withhold either every fifth or every thirteenth point from the model-fitting analysis.

In taking the daily averages, the temperature was averaged over the time the fan was on, rather than over the entire 24-hour period. The weighting factor established for use with the daily averages was the count of power-on records making up the average.

Deriving and Checking the Models

The models were fit based on a least-squares estimation of the parameters. The models were:

$$\text{power} = \left\{ \begin{array}{c} A + B \cdot T_{out}, \\ A + B \cdot T_{out}^2, \\ A + B \cdot T_{out} + C \cdot T_{out}^2, \text{ and} \\ A + B(T_{out} - T_{break}) \cdot (T_{out} > T_{break}) \end{array} \right\}$$

where A, B, C, and T_{break} are the parameters to be estimated. Since the logical expression in the last model evaluates to 1 or 0, continuity is ensured at the breakpoint.

Each model was estimated using the hourly data and the daily average information. Additionally, the daily average data were weighted by treating each power-temperature record as if it had been duplicated by the number of hours that had been averaged into the record—that is, by the number of hours the fan was on that day. With the best-fit models in hand, the data held back from analysis were used to check the results. Just as the models were fit using the least-squares criterion, they were checked by summing the squares of the differences between the actual and predicted powers for each record. To provide a basis of comparison between fans, the sum of squares of the independent data about their mean was recorded as well. A typical summary chart appears as Table 2.

This table indicates that for the 140 records held back from the analysis of supply fan 10 operating under variable-inlet-vane control, the best-fit model was the piecewise linear one derived using the hourly data. The small error sum of squares entry indicates that the 146 kW^2 total variation about the mean was explained most effectively by the power-temperature relationship expressed in that model.

Not all the model comparisons were as straightforward as this, however, and in several cases an analysis of variance was calculated in order to determine whether the observed differences between models and between hourly, daily, and weighted daily averages were meaningful at some significant level or could be explained just as easily by chance. Often the results depended on the subgroups considered. For example, referring to the error sums of squares presented in Table 2,

- the apparent differences between using the hourly, daily, and weighted daily averages cannot be said to exist at any standard level of significance;

TABLE 2
Summary Chart for Comparison of AH-10-SF Models for VIV Operation

Predict 140 records, mean=10.430, ss about mean=146.00	Error sum of squares		
SF10 VIV model	hourly data	daily averages	weighted daily
Linear regression	79.24	79.09	79.46
Quadratic-1	60.03	59.95	60.32
Quadratic-2	35.87	36.30	36.64
T_{break}	31.05	32.26	32.66

- considering only the T_{break} and quadratic-2 models, the apparent differences between them may be said to exist at the 25% but not at the 10% level of significance; and
- the apparent differences between the daily and weighted daily averages alone may be said to exist at the 0.1% level of significance or better.

In other words, for the models summarized in Table 2 for supply fan 10 under VIV control, there is no reason to prefer those based on hourly averages to those based on daily averages, although the daily average models almost certainly are better than the weighted daily average models. The piecewise linear model is better than the quadratic model, with a probability of between 75% and 90%.

In general, we came to expect that the piecewise linear model would provide the best fit, the simple linear regression would provide the poorest, and the hourly averages would yield better fits than the daily averages but not by any statistically significant margin. Weighted daily averages generally were found to provide closer estimates than daily averages (even though this was not true for the data in Table 2).

In only one case—that of AH-8-SF—was the piecewise linear fit found to be inferior to one of the other models. For VSD control over this fan, the three-term quadratic model showed slightly smaller error sums of squares. However, the best-fit estimation contained a vertex at 30°F; this vertex, below which the model shows fan power rising as temperature decreases, suggests that a piecewise quadratic model be fit to the data. For consistency with the other three fans, however, the piecewise linear model was used instead. A graph of the data and the chosen line for AH-8-SF appear as Figure 1. This figure also gives an indication of the degree of scatter inherent in the power data when considered only as a function of outside temperature.

Comparing the Models

After power vs. temperature models had been developed for all four fans, the error sums of squares from charts similar to Table 2 were compared for each fan using analysis of variance techniques. To do this, the error sums of squares first were normalized by dividing by the appropriate sum of squares about the mean. The resulting statistic is 1 minus R-square, the familiar test of goodness of fit for a linear regression. Values ranged from a worst case of 1 (for two-parameter fits to daily averages on AH-11-SF) to a best case of 0.2 (for three-parameter fits to AH-10-SF). One might expect that the degree of variability in the performance of a fan has the greatest effect on the goodness of fit any model might achieve, and indeed, the effect of fans was found to be significant at the 0.1% level or better.

Besides testing for differences attributable to the fan, we checked to determine whether there was any basis for stating that the piecewise linear model was the best overall and whether the apparent distinctions between hourly, daily, and weighted daily models were supported by the data. Taken all together, the values for R-square provide no basis for assuming that any of these observed differences were anything but chance occurrences.

This last result indicates that daily averages are just as effective as hourly averages when modeling power as a function of outside temperature. Clearly, this has broad implications for simplifying the data collection. However, our sense is that with only four fans in the data base, the observed differences are slight enough not to be significant, but that if more fans were analyzed in a like fashion, a difference could be detected at some reasonable level of significance. Therefore, we would say that the hourly averages outperform the weighted daily averages, and the weighted daily averages outperform the daily averages, but that given an eight-month data collection effort, the differences are slight and may not justify the added costs of collecting and storing data on an hourly basis.

In every analysis, the model was found to be significant at some standard level. With all the models taken together, the differences found between them were statistically significant at the 0.1% level or better, and whenever the piecewise linear model was compared to the other models, or to the three-term parabolic model only, it was

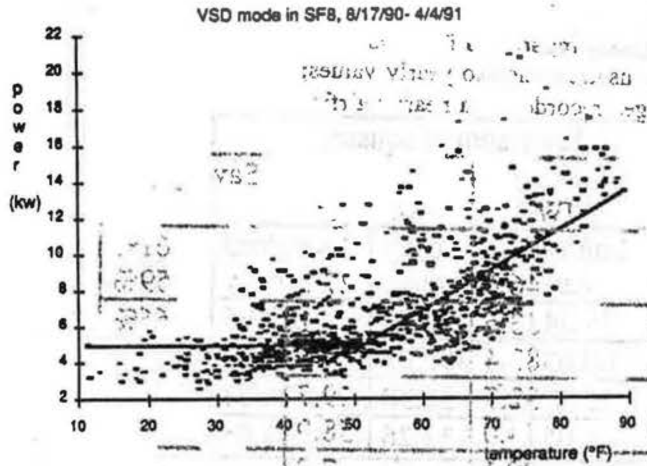

Figure 1 *These data, for VSD mode AH-8-SF, were the only set for which the piecewise linear fit could not be distinguished as the best of the two- and three-parameter models tried. As shown here, the piecewise linear model was chosen to provide consistency with the rest of the analyses.*

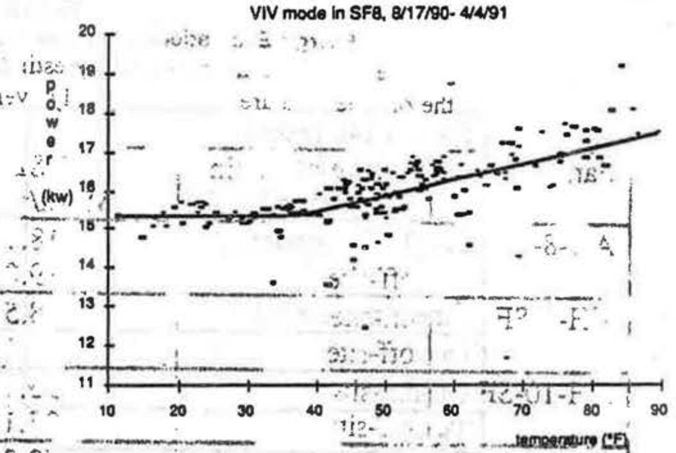

Figure 2 *The variable-inlet-vane mode for AH-8-SF. The total data base for this fan consisted of 261 hours of operation.*

significant at most at the 10% level and often at the 1% level. Thus, one may state that the piecewise linear model is better than its two- and three-parameter competitors, with at most a 10% chance of error.

Several representative plots of power against outside temperature, along with the best-fit models determined, are shown as Figures 1 through 3.

FAN ENERGY CONSUMPTION

Experimental Results

The empirical correlations of fan power with outside temperature were applied to two sets of binned data: temperatures recorded on-site over the eight-month period and long-term yearly average data from a nearby airfield. Both were binned in 5°F intervals, and both assumed that in each bin, the number of hours the fan was turned on could be taken as the bin count times the overall fraction of time the fan was on. This fraction was found by taking the full number of points the fan was on plus half the number of points filtered out as mixed-power points and dividing by the total number of records originally collected.

Because nighttime hours, during which the fans are turned off, are cooler on average than daytime hours, assigning on- and off-hours this way will overestimate the number of hours the fan was on at lower temperatures. Therefore, the energy estimates based on binned data were expected to be underestimates, and indeed they tended to be 1% to 4% lower than numbers found using the actual temperatures along with the knowledge of when the fan was on and off.

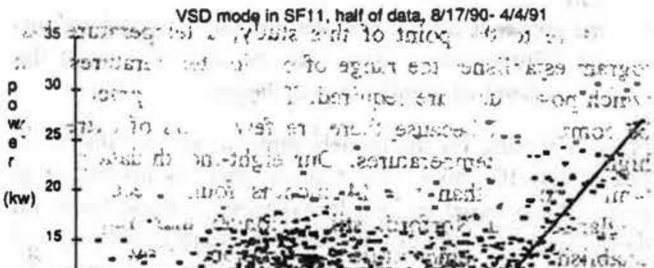

Figure 3 *The variable-speed-drive mode for AH-11-SF. This fan showed the highest breakpoint temperature of all, at 71°F.*

Table 3 shows the results from the eight-month survey period extrapolated to one year and indicates substantial, and variable, energy savings.

The close agreement between these estimates and those obtained using long-term average data indicate that the eight-month monitoring period accurately represents the distribution of temperature data over a year. This agreement suggests that another method for estimating annual savings would be to estimate the minimum time period required to emulate the annual temperature distribution and record fan power over that period without attempting to establish correlations with temperature. In fact, for VSD control, the average of the collected data for each fan was found to supply a least-squares estimation of *total* energy use in many cases as good as that derived using the best-fit model (although of course the average was not a good predictor of individual power records). Intuitively this makes sense, as the fans ran in the VSD mode for the bulk of the eight months during which the data were collected.

TABLE 3

Energy Estimations Based on Binned Temperature Records

(The eight-month on-site energy estimations are scaled to yearly values; the off-site data are long-term yearly averages recorded at a nearby airfield.)

Fan	Temperature bin source	VSD MWh/yr	VIV MWh/yr	Savings, MWh/yr	Savings, %
AH-8-SF	on-site	18.1	46.5	28.4	61%
	off-site	19.0	46.6	27.6	59%
AH-9-SF	on-site	28.5	83.6	55.1	66%
	off-site	29.1	83.4	54.3	65%
AH-10-SF	on-site	33.1	92.7	59.7	64%
	off-site	35.1	94.0	58.9	63%
AH-11-SF	on-site	52.3	100.1	47.9	48%
	off-site	54.0	100.6	46.6	46%
Total	on-site	132.0	322.9	191.1	59%
	off-site	137.2	324.6	187.4	58%

More to the point of this study, a temperature histogram establishes the range of outside temperatures over which power data are required. This range in practice can be compressed because there are few hours of extremely high and low temperatures. Our eight-month data period, while shorter than the 14 months found necessary by Englander and Norford, still is longer than required to establish the temperature correlation. However, our experience with weighting the data indicates that, should an attempt be made to draw out a full correlation using only data collected over representative shorter intervals, the experimenter should ensure that the distribution of observed temperatures matches that of the expected yearly temperature histogram.

A Priori Savings Estimates Based on Manufacturers' Specifications

Motivation and support for this research stem from the uncertain and unproved savings estimates made without benefit of measurements. The annual savings estimates above, grounded in data, now should be compared with what was expected. As a first step in this comparison, we present the savings estimated by easy-to-use software supplied by two VSD manufacturers. This software requires a flow histogram, hours of fan operation, and motor size. The two manufacturers employ different default flow histograms, the first narrowly distributed about 70% of full flow and the second a broader distribution centered at 55% of full flow. Figure 4 compares the flow histograms used by these manufacturers.

Using these flow histograms, savings estimates now can be made. Table 4 shows the percent savings that depend only on the flow histogram. Savings are higher with the broad, low flow distribution, since prolonged operation at high flows causes the variable-speed drive to run the fan speed up to that used under VIV control.

We note that the correlation between fan power and flow incorporated in the first set of software is somewhat crude. For inlet vane control, fan power is assumed to vary linearly with flow. For the VSD mode, the classic cubic relationship is used, even though static pressure control, typically implemented for supply fans, will yield less variation of power with flow. Data reported by Englander and Norford show that both inlet vane and VSD control have a strong quadratic component. The power-airflow correlations used by the second manufacturer appear to match the reported data more closely. However, differences in the manufacturers' power-airflow correlations do not appear to influence the estimated savings. Significantly, the savings depend strongly on airflows. For the default distribution used by the second manufacturer, the percent savings match those achieved by the VSDs we have monitored.

Absolute savings, on which finances are based, are calculated from two pieces of information besides the flow

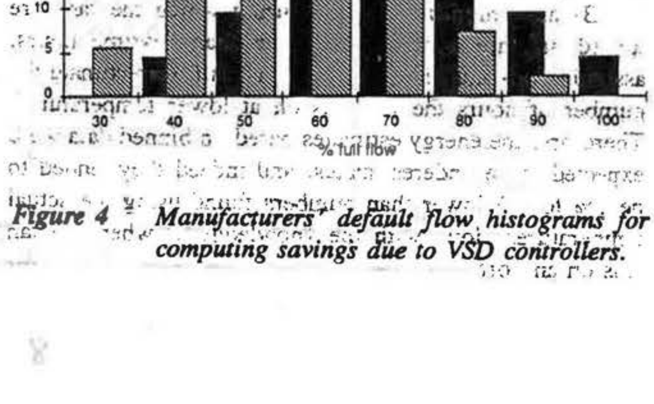

Figure 4 Manufacturers' default flow histograms for computing savings due to VSD controllers.

TABLE 4
Manufacturers' Computed Percent Savings Depend Only on the Flow Histogram, Not on Hours of Operation or on Motor Size

Manufacturer	Savings, high flow distribution	Savings, low flow distribution
1	42%	58%
2	41%	59%

histogram: hours of operation and peak motor power. Hours of fan operation can be obtained from the building operator or, in this case, from direct measurement. Fractional runtimes for the four fans ranged from 0.33 to 1.00; the variation indicates that *a priori* assumptions about operating hours would not lead to accurate savings estimates. Table 5 shows that fan motor nameplate power typically overestimates the actual power required to supply design airflow. If, instead, the brake horsepower—at which the fan curves indicate the system provides rated airflow—is used to calculate the absolute savings, a more accurate figure results.

It appears that the broader flow histogram centered at lower flow values is a reasonable choice for these four fans. Further, it is essential to specify two important parameters accurately: hours of operation and the fan power required at design conditions.

CONCLUSION

An average energy savings of 58% was achieved for four supply fans for which inlet vane control was replaced with variable-speed motor drives. Savings for individual fans ranged from 46% to 66%. These savings were calculated based on best-fit models in the least-squares sense. Of the two- and three-parameter models considered, the piecewise linear model provided the closest fits. In general, the use of hourly data is encouraged when developing these fits, and if daily average powers are used, they should be weighted for hours of operation.

The aggregate savings for these fans can be reproduced closely by using VSD manufacturers' software with a broad flow histogram centered at 55% of full flow and accurate estimates of operating hours and fan power at design airflows. However, it must be remembered that all four fans studied were designed at the same time and installed in the same building. Moreover, the flow distribution typically is not known *a priori* and the agreement between estimated and measured savings is not guaranteed in general. In addition, the manufacturers' calculations do not account for changes in static pressure setpoint for the supply fans, a parameter that has previously been shown to have a strong influence on supply fan power.

End-use monitoring must be carefully integrated into conservation assessment programs. When such monitoring is deemed necessary, outside temperature data offer an attractive alternative to airflow measurements as a means of partially explaining variations in fan power.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the EUA Cogenex Corporation, MIT Physical Plant, Northeast Utilities, Texas A&M University, and the Empire State Electric Energy Research Corporation.

TABLE 5
Comparison of Savings Estimated by Empirical Fits to Outside Temperature and VSD Manufacturers' Software, Using Default Flow Histogram, Hours of Operation, and Peak Fan Power
(The motor nameplate horsepower overestimates total savings by 42%, while taking brake horsepower from Table 1 overestimates by only 6%.)

Fan	Savings, MWh/yr		
	Empirical fit to outside temperature	Flow histogram with nameplate horsepower	Flow histogram with brake horsepower from Table 1
AH-8-SF	28.4	24.4	20.5
AH-9-SF	55.1	70.6	45.9
AH-10-SF	59.7	117.1	82.0
AH-11-SF	47.9	59.6	53.6
Total	191.1	271.7	202.0

REFERENCES

Box, G.E.P., W.G. Hunter, and J.S. Hunter. 1978. *Statistics for experimenters*. New York: John Wiley & Sons.

Englander, S.L., and L.K. Norford. 1992a. Saving fan energy in VAV systems, part 1: Analysis of a variable-speed-drive retrofit. *ASHRAE Transactions* 98(1).

Englander, S.L., and L.K. Norford. 1992b. Saving fan energy in VAV systems, part 2: Supply fan control for static pressurization minimization using DDC zone feedback. *ASHRAE Transactions* 98(1).

Fels, M.F. 1986. PRISM: An introduction. *Energy and Buildings* 9(1&2).