Commercial Air-Conditioner Load Control Results From Multiple Programs

Dr. Michael Perry, FSC Group, San Francisco, CA

ABSTRACT

Air-conditioner (AC) load control programs for commercial customers have several benefits as a complement to programs for residential customers. Commercial AC load peaks at a different time of day than residential load, which provides more flexibility in providing demand response at different times. Also, in many cases one business owner may have the authority to sign up many AC units for load control, which can make recruiting more efficient. Finally, commercial AC load is less weather sensitive than residential, which means that commercial AC load control programs can provide substantial demand response even at only moderately hot temperatures.

Although AC load control is potentially one of the simplest types of demand response, it has several complexities when used for commercial customers: recruiting businesses for the program can be difficult for multiple reasons. Commercial AC load magnitudes vary more across customers than do residential loads, which implies the need for larger sample sizes to precisely estimate commercial AC load impacts. The same issue also increases the importance of targeting customers likely to produce large load impacts.

In the introduction to this paper the primary details of commercial AC load control programs are introduced and illustrated using results from evaluations of four AC load control programs for the summer of 2010. The remaining sections of the paper are used to discuss the above complexities of these programs and their implications.

Introduction

This section provides background on commercial AC load control programs and basic facts and results from four recent evaluations of such programs. It also has a discussion of the benefits of commercial AC load control programs as a complement to residential programs.

Commercial air-conditioning (AC) load control programs are demand response resources in which commercial customers agree to have control devices installed on their AC units that allow the utility to remotely turn down their AC at times of peak demand. These programs typically operate as emergency demand reduction resources, but sometimes also serve as energy-saving resources.

The periods during which the utility controls the AC units are called events and usually last two to six hours. The control devices are either load-control switches (referred to here as switches) or programmable communicating thermostats (PCTs).

Switches are installed on the AC unit, usually exterior to the building, and are programmed to directly control the AC duty cycle. Control strategies currently used for switches are either simple cycling or more complex algorithmic cycling. Under simple cycling, a maximum duty cycle is imposed on the AC unit—for example, 50% duty cycle. If a given AC unit would have been running above 50% during the event, then the effect of the event is to reduce the duty cycle to 50%. If the AC would have been running below 50% during the event, then the event has no effect on AC load. Under algorithmic cycling, the switch has some ability to learn about an individual AC unit's typical duty cycle during hot days. On these learning days, the switch records hourly AC load to use to predict usage on other hot days. During an event, the switch controls the AC unit so that it uses a pre-determined fraction of the energy it would have normally used under the same conditions. For example, a 50% algorithmic strategy would reduce an AC running at a 50% duty cycle to a

25% duty cycle. If the switch does not have enough chance to learn about the AC unit's typical duty cycle before an event, then the switch defaults to simple cycling. For example, a 50% algorithmic switch with too few learning days would default to 50% simple cycling.

PCTs are installed inside the premise, and act as normal programmable thermostats, except during event periods. During events, the utility sends a signal to the PCT, which either tells the PCT to set the temperature higher (temperature setback) or tells the PCT to control the duty cycle of the AC in the same way a switch would. A typical setback strategy for a four-hour event would be to set the temperature two degrees higher in the first hour and then one additional degree higher in each of the next two hours. The benefit of a temperature setback strategy over a cycling strategy is that it provides load impact for any level of AC duty cycle and it requires no learning days. The downsides of using a PCT are that the unit itself is more expensive than a switch, it requires installation inside the premise which means an appointment has to be made, and the fact that it is inside makes the PCT less reliable at receiving the control signals because it can be blocked by walls.

Four commercial AC program results are introduced here and used in subsequent sections for illustration: Pacific Gas & Electric's (PG&E's) SmartAC program; Southern California Edison's (SCE's) Summer Discount Plan; San Diego Gas & Electric's (SDG&E's) Summer Saver program; and Ontario Power Authority's (OPA's) *peaksaver* program. The programs will be referred to by the name of the utility rather than the program name. The results discussed here are documented in detail in the four reports listed in the References section.

Two of the programs—those at OPA and SDG&E are targeted exclusively towards small and medium businesses (SMB—roughly defined as those with less than 200 kW peak demand). PG&E's program is also comprised almost exclusively of small and medium businesses. SCE's program is targeted at large, medium and small commercial customers.

Table 1 shows basic information about the four programs as of summer 2010. The first column shows the total number of commercial participants, which ranges from about 2,500 to about 10,000. SCE and SDG&E's programs each have no PCTs, using switches exclusively, while PG&E's and OPA's programs each have a substantial number of PCTs.

Event counts, durations and temperatures in the table include both population-wide events and test events directed only at a sample of customers for the sake of evaluation.

Table 1. Summary Information for Four Commercial AC Load Control Programs

Utility	Participating Customers	Percent of Commercial Tonnage that is SMB	Percent PCTs	Average Event Duration
PG&E	2,446	99	86	4 hours
SCE	9,665	67	0	2 hours
SDG&E	6,307	100	0	4 hours
OPA	2,900[1]	100	33	4 hours

Control strategies vary widely across programs, as shown in Table 2. The setback descriptions in the table refer to the number of degrees Fahrenheit that the PCT is adjusted upward during each hour of the event. For example, 2-1-1 refers to a strategy of raising the temperature 2°F in the first hour and 1°F in each

[1] Estimated from the number of control devices, 4,430. Actual number of participating customers is not available.

of the next two hours, then leaving it the same for the rest of the event. A setback strategy of 4 refers to a 4°F temperature raise in the first hour that lasts the entire event.

Table 2. Control Strategies

Utility	Switch	PCT
PG&E	30% algorithmic cycling	2-1-1 setback
SCE	30%, 40%, 50%, 100% simple cycling	NA
SDG&E	30%, 50% simple cycling	NA
OPA	30%, 50% simple cycling	1-1-1-1 setback or 4 setback

For each program, a sample of customers had AC loggers installed for the summer of 2010. The loggers recorded AC usage data at 5-minute intervals for the whole summer. This data was analyzed at the hourly level. Reference loads and load impacts were estimated for each program using regression analysis. Regression analysis was performed separately for each program and consisted of an individual linear regression for each customer in each sample with variables that controlled for weather, time of day, day of the week, month and the presence of an event. Regression specifications differed across the four programs, but were constant across customers within a program. In the cases of PG&E and SDG&E, regression analysis was corroborated by a treatment-control analysis, discussed below.

Table 3 shows the size of each AC logger sample and the results from each analysis. Reference loads vary widely across the programs, from a low of 0.44 kW at SDG&E to a high of 2.09 kW at OPA. Load impacts also vary widely, both in absolute magnitude and as a fraction of reference load. These differences are due to several factors. Most importantly, very high load impacts at SCE can be attributed to very hot weather, which leads to high duty cycles and therefore large load impacts due to cycling. At PG&E, low impacts were attributed to substantial communication failure, in which control devices do not receive the signal for an event. Different control strategies also have an effect on the load impacts, but it is not currently clear exactly which strategies work best under which conditions.

Table 3. AC Logger Sample Sizes and Estimation Results

Utility	Sample Size (AC units)	Number of Events	Event Duration (hours)	Average Event Temperature (°F)	Reference Load (kW/AC unit)	Load Impact (kW/AC unit)	Percent Load Impact
PG&E	300	17	4	86	1.54	0.11	7
SCE	400	2	2	101	1.90	0.90	47
SDG&E	500	11	4	84	0.44	0.09	21
OPA	400	6	4	86	2.09	0.56	27

For each utility, the total number of commercial participants is much smaller than the number of residential participants, although the total size of each program's residential program varies widely from about 25,000 customers in SDG&E's program to about 330,000 customers in SCE's program. This is typical. In North America, AC load control programs typically contain a residential branch that provides the bulk of load impacts and a commercial branch that contains many fewer customers and accounts for a small fraction of the total load impact. This is true despite the fact that commercial AC units themselves generally have higher loads and load impacts than residential units, and that commercial customers often have multiple AC units to control, while residential customers usually have only one.

Commercial customers have substantial potential as cost-effective participants in these programs for three reasons. First, despite some complications in recruiting (discussed in the next section), commercial customers often have many AC units that can be signed up for the program all at once. This can lead to reduced marketing and recruiting costs.

Second, commercial AC load peaks earlier in the day than residential AC load. Commercial loads typically peak around 4-5 PM on weekdays, while residential loads peak around 6-7 PM. This means that commercial customers complement residential customers by providing larger load impact at times when residential customers provide smaller impacts. Suppose a utility is facing a potential peak day and would like to rely on AC load control to moderate the peak. It may not be clear exactly when the peak will occur. Having a large commercial AC load control program in addition to a residential program implies that the demand response potential of AC load control could be substantial over several hours of the day rather than just in the evening.

Finally, commercial AC load is less weather-sensitive than residential AC load, providing the potential for substantial load impact at times of less extreme temperatures.

With these useful aspects of commercial AC load control in mind, the remainder of this paper discusses some complications that need to be considered and addressed for a commercial AC load control program to be cost-effective as either an emergency demand resource or as an energy conservation resource.

Recruiting Commercial Customers into AC Load Control Programs

The full range of reasons that commercial AC load control programs remain smaller than residential programs is not explored here. However, it is clear that getting commercial customers to sign up for AC load control is challenging.

There are several reasons that it can be difficult to attract commercial customers. First, commercial AC load control marketing is often done through direct mail, which typically yields a sign-up rate of only 1-2%. Telemarketing can give much higher sign-up rates, but at a higher cost. A recent telemarketing effort for a commercial AC load control program done by Population Research Systems (PRS, a component of the FSC Group) suggests two major improvements to program marketing. In this effort, PRS achieved a response rate of almost 10%. However, there was significant difficulty in reaching the correct decision-maker at each business and many businesses were not eligible for the program because they did not have AC, they had moved out, or the provided contact information was invalid. Of eligible businesses, with correct contact information and where the correct decision-maker could be reached, the sign-up rate was 24%. The difference between the overall sign-up rate and the rate for decision-makers reached, suggests that marketing could be substantially improved by making sure that contact information is correct, by targeting businesses likely to have AC and by finding ways to get through to the decision-maker at a business. These points also are likely to improve direct mail campaigns.

A second reason that it can be difficult to attract commercial customers is that business managers may be wary of the risk of making customers uncomfortable on hot days. While this is a valid concern, business managers may over-estimate the degree to which AC load control leads to discomfort, particularly given the generally mild control strategies used by many programs. During the summer of 2010, PRS conducted surveys about discomfort during events for both PG&E and OPA commercial customers. In each case, the difference in the fraction of respondents who reported employee or customer discomfort during a load control event and the fraction of respondents who reported employee or customer discomfort during a control period was only about 10-12% of the sample. That is, only about 10-12% of the samples experienced more discomfort than they would have had there been no event at all. It could be useful in program marketing to make this known to customers.

Commercial AC Load Impact Estimation

The primary benefits of a commercial AC load control program are likely to be avoided generation capacity costs. A program will also provide a smaller amount of avoided transmission and distribution costs and may also provide energy savings. The costs of the program are primarily those associated with recruiting and compensating customers, program administration and maintenance and measurement and evaluation costs.

In order to determine whether a program is cost effective, load impacts need to be estimated fairly precisely. Estimating load impacts requires gathering AC load data on a sample of customers and observing loads during load control events. This is currently done primarily using AC loggers, which record load data from the AC at 1-minute or 5-minute intervals. Installing AC loggers is expensive and sample sizes are limited by this expense. As more customers have interval meters installed, impact estimation for residential customers will be done primarily using meter data. Whether this can be done for commercial customers reliably is still an open question and is discussed at the end of this section.

Estimating load impacts using AC logger data for a sample of program participants can be done using two main modeling strategies: first, a model of what AC load would have been but-for an event can be estimated using AC loads for the same set of participants on non-event days with similar conditions; second, a model of AC load but-for an event can be estimated using AC loads for similar AC units not subject to load control at the time of the event. Using the first strategy requires a model that predicts AC load at a given time as a function of observable variables. This usually takes the form of a regression model or a day-matching model. The variables available for modeling a given customer's AC load are typically the date and time, and the weather that the customer was exposed to at that time.

Using the second strategy requires the ability to observe both a treatment group of AC units that will experience an event and a control group—ideally containing AC units that are in the program—that does not experience an event. Implementing this strategy therefore requires that the AC load control provider has the ability to control subsets of AC units within the program. On an event day, the provider only implements the event on one subset of the program participants, while another subset is not controlled. The group of customers not controlled then provides information about what load would have been in the group that was controlled. If this is feasible, this treatment-control strategy has the strong advantages over the first strategy that it does not require an assumed model of AC load as a function of weather and it does not require the existence of non-event days with hot temperatures to use for modeling.

In addition to measuring impacts themselves, most evaluations also require a measurement of the variation in impacts expected under a given set of weather conditions.[2] Variation in impacts for a given set of conditions arises from at least three major sources. First, participant behavior and AC unit load varies for many reasons other than the immediate weather conditions. This source of variation is made up of two components, an observable component and an unobservable component. The observable component is observable because differences in AC load at similar conditions during a given summer are recorded in the data collection process. This component consists of all the variation in AC that occurs and is observed, but that is not captured by a weather-based model of AC usage. Examples include participant vacation timing or a business' summer schedule. The unobservable component consists of the changes in AC load that will take place in the future but that are not represented in the data recorded for a given evaluation. This unobservable component consists of things like changes in the economy or changes in AC efficiency.

[2] These measures are not reported here, but are reported in the evaluations themselves, which are listed in the references.

Second, AC load measurements made using a sample of customers can vary from the values in another sample or in the full population due to random differences between the sample and the population. These differences can also be broken into an observable piece and an unobservable piece. An example of the observable piece would be if one sample contained a larger proportion of customers from one particular region than another sample. Weighting of results can be used to account for this variation. An example of the unobservable component would be if one sample just happened to contain more high use customers than another even though all the observable variables for the customers are the same. This type of variation generally cannot be accounted for accurately, but will be larger if sample sizes are smaller or if customer load shapes or magnitudes vary more across customers.

Third, the program itself can change. Different types of customers can sign up. The utility may change aspects of the control strategy or the timing of events.

The observable components of the first and second sources are usually explicitly modeled in evaluations of these programs. The unobservable components of the first and second sources of variation and the third source are not usually modeled, but are implicitly understood to be present as well and to provide additional uncertainty about what impacts to expect from the program.

There is an important difference between commercial AC loads and residential AC loads that is relevant to the unobservable part of the second source of variation: commercial AC loads vary more in magnitude across AC units and across customers than do residential AC loads. This means that the unobservable component of the second source of variation will be larger for commercial customers than for residential customers. Again, this variation usually cannot be accurately removed or even measured, but must be understood to be present.

Figure 1 shows hourly average AC loads for a sample of 300 residential AC units and Figure 2 shows the same for a sample of 300 commercial AC units from the same AC load control program. This data was collected using AC loggers during the summer of 2010. Each curve shows the hourly average load for one decile of AC Units, where deciles are calculated based on average AC load on days when the temperature exceeded 85°F. By comparing the two figures, it is clear that there is a much wider distribution of AC loads for commercial AC units than for residential. The highest decile of residential AC units has an average peak load of just over 2 kW, while for the commercial units the highest decile has an average peak of almost 5 kW. These differences in size are also likely to be correlated with differences in average AC load shape over time, leading different sized air-conditioners to achieve peak load at different times of day.

Greater heterogeneity in load means that a given sample of AC units is less likely to accurately reflect the full population of program participants, and this applies regardless of which modeling strategy is used. Therefore, for the same number of AC units sampled, there will be more uncertainty in commercial AC load impact estimates than in residential AC load impact estimates. Alternatively, this implies that it is more expensive to estimate AC load impacts for commercial AC units to the same degree of precision as impacts for residential AC units because it requires a larger sample.

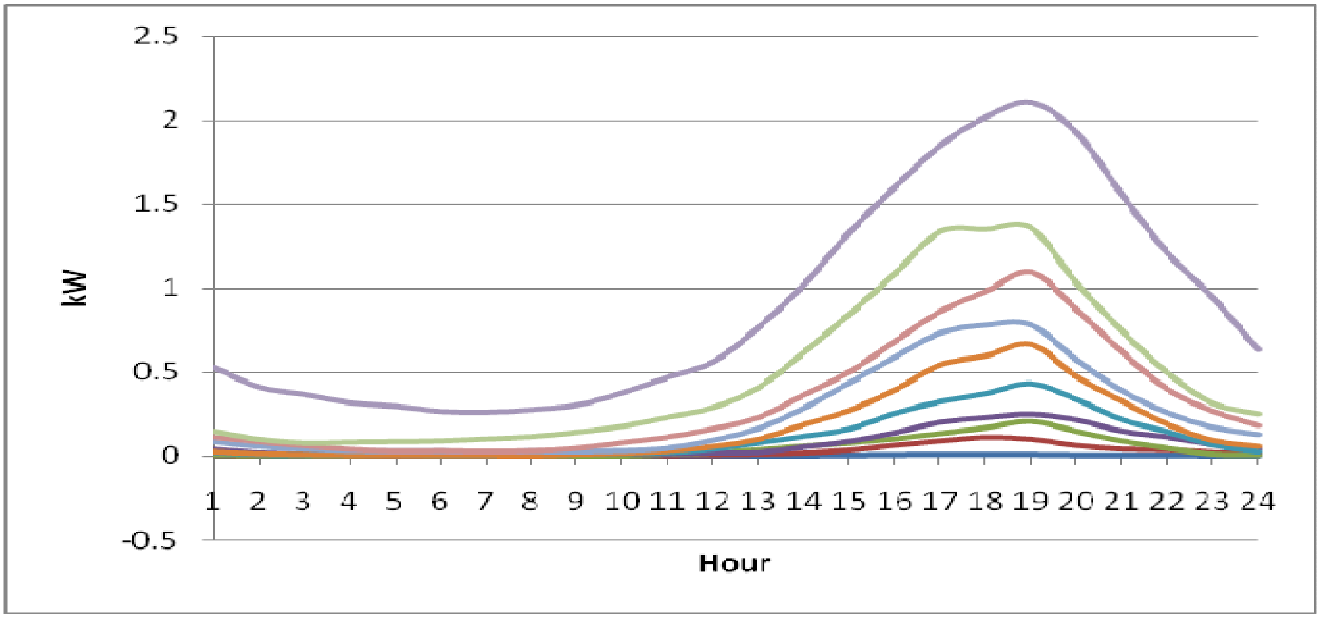

Figure 1. Average Hourly AC Loads For Residential AC Units, by Decile of Average AC Load

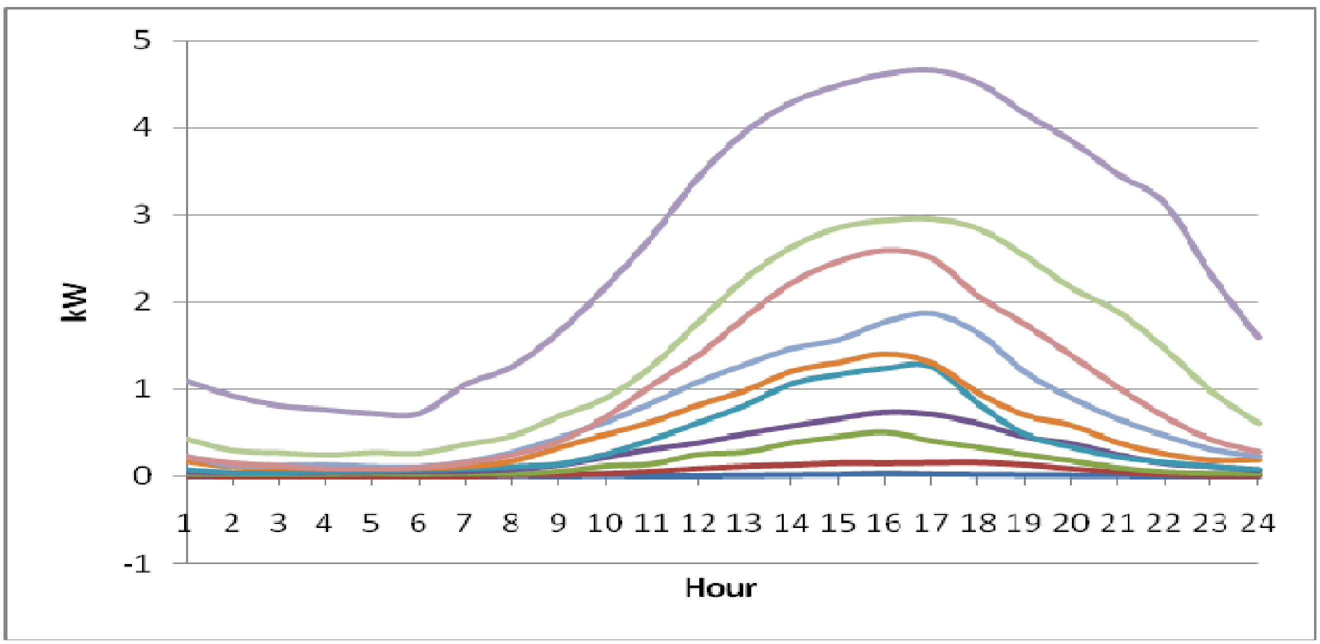

Figure 2. Average Hourly AC Loads For Commercial AC Units, by Decile of Average AC Load

The same issue also has the implication for the treatment-control modeling strategy that it requires larger samples for the control group to accurately represent the behavior of the treatment group if there had not been an event. An example of this is shown in Figures 3 and 4. Figure 3 shows AC load for a treatment and control group each consisting of 115 residential AC units and Figure 4 shows AC load for a treatment and control group each consisting of 220 commercial AC units. Both figures are from the same AC load control program and the same day of the summer of 2010. In each case, the treatment and control group were drawn randomly from the sample of customers that had AC loggers installed, meaning that there should be no systematic differences between the two groups. The effect of the event on the treatment group is evident in both figures starting at 2 PM (hour 15 in the figure). The important aspect of the two figures is that during the hours before the event, the treatment and control groups in the residential sample match each other much better than the treatment and control groups in the commercial sample. In this case, the difference between the two curves in Figure 4 prior to the event strongly suggests that event impacts calculated by subtracting the load of the treatment group from the load of the control group would lead to a significant underestimate of the event impact. In contrast, in Figure 3, it appears that such a strategy would lead to an accurate estimate of the event impact. This is despite the fact that the commercial sample is almost twice as large as the residential sample. With a wider distribution of average AC load in the population, the variance of average load in a randomly-drawn group of a given size is larger, and therefore it is more likely for the two groups to not match each other well.

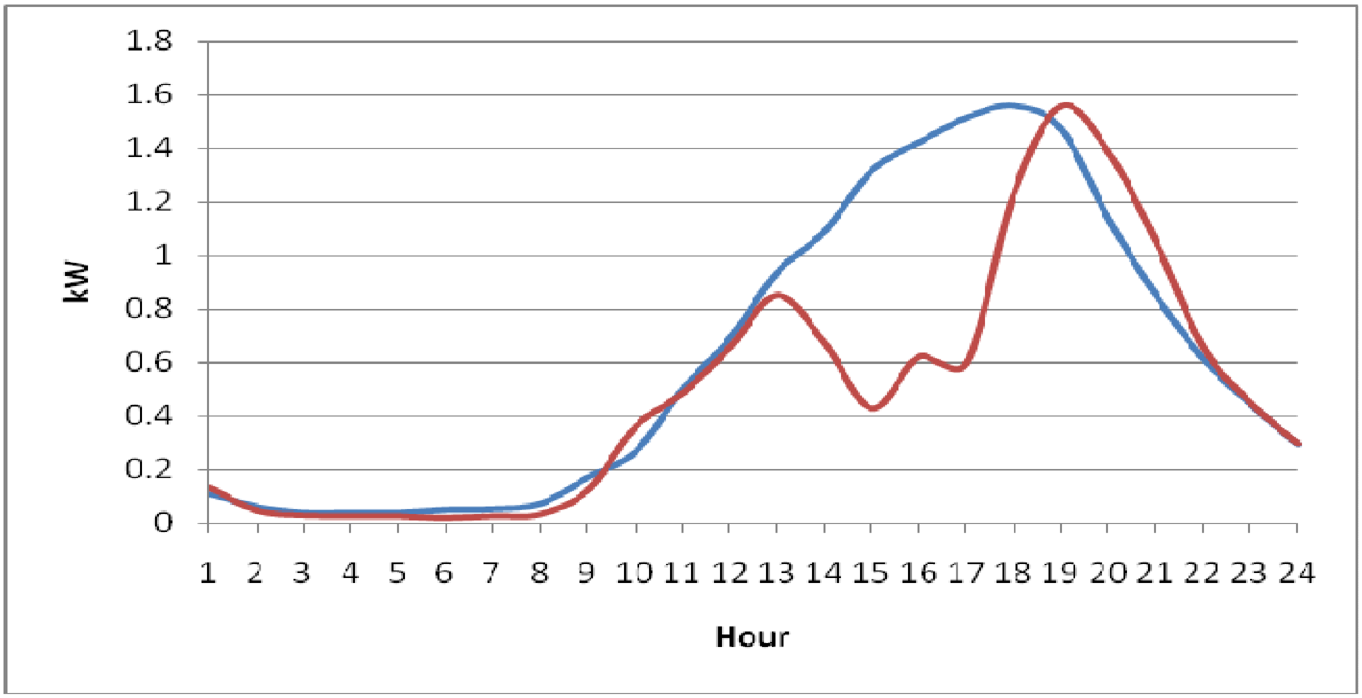

Figure 3. Average Hourly AC Loads For A Residential Treatment and Control Group

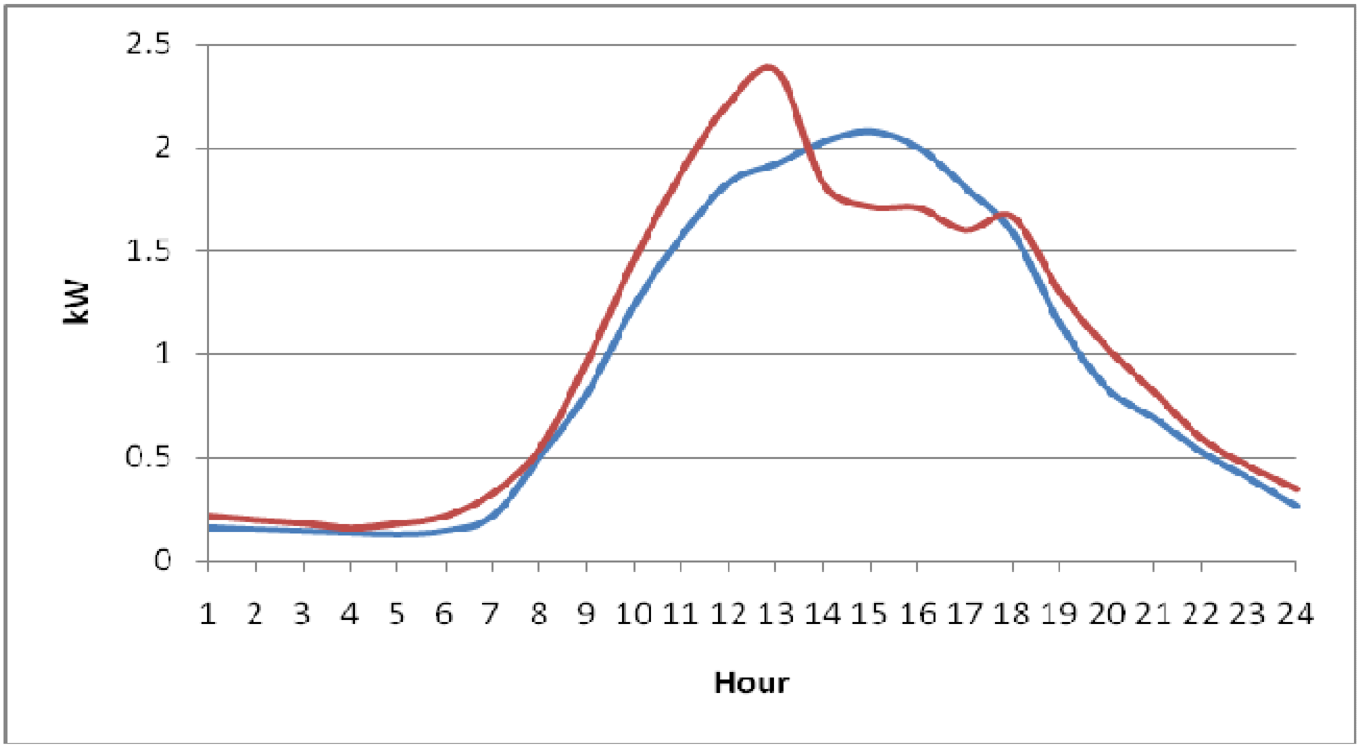

Figure 4. Average Hourly AC Loads For A Commercial Treatment and Control Group

There is no analytical way to eliminate this issue. The only thing to be done is to recognize that sample sizes need to be larger for commercial customers to achieve the same level of precision in load impact estimates as for residential customers. Alternatively, it must be expected that the impacts from a commercial program under a given set of weather conditions will vary more than those from a residential program.

As mentioned, residential load impact estimation will mainly be done using whole-building interval data in the future as most customers have interval meters installed. Load impact estimation using whole-building interval data works well for residential customers because the reduction in AC load due to an event is uncorrelated with other changes in household load and because AC load is a substantial part of residential whole-building load. Load impact estimation using whole-building interval data for commercial customers is complicated by the fact that AC load is often a smaller fraction of total load and by the fact that commercial customers' load impacts are often smaller because load control strategies used on commercial customers are often quite mild. Both of these issues mean that the effect of the event is harder to accurately estimate against the background noise of the whole-building load. With large enough sample sizes, it should be feasible to estimate commercial load impacts with whole-building data, but it will require greater computing resources (for running regressions on many customers' worth of data) or it will require the willingness to use quite large treatment and control groups in order to precisely measure event impacts against the underlying noise.

Customer Targeting

Commercial customers can be targeted for AC load control to a greater degree than residential customers due to the fact that particular industries systematically have larger AC loads and therefore provide larger load impacts during events. Figures 5 and 6 show average hourly per AC unit load for two of the programs in 2010. Despite the fact that industry classifications are not consistent across programs, a couple noticeable patterns emerge. First, retail stores and offices typically have larger average loads. Second, non-commercial enterprises, such as churches, schools and government offices have quite low average loads.

Ideally, a load-impact based model of cost-effectiveness would be developed for a given program, and then customers would be profiled based on industry, building size, geographic region and any other relevant information in order to provide lists of customers that are likely to provide the cost-effective level of load impact. Although this would entail upfront costs for data gathering, it would likely pay off substantially by allowing programs to avoid targeting customers unlikely to provide substantial load impacts.

Although it is most cost-effective to target customers with substantial AC loads, it is also the case that those customers may often be the ones least likely to sign up for AC load control programs. A business paying for a large amount of air-conditioning may be doing so because the owner perceives climate control as important for profitability. For example, a restaurant or movie theater may have substantial AC load, but also may be unlikely to risk making customers uncomfortable. The best participants to target may be those with substantial load, but where customers do not stay very long, such as a convenience store or fast food restaurant.

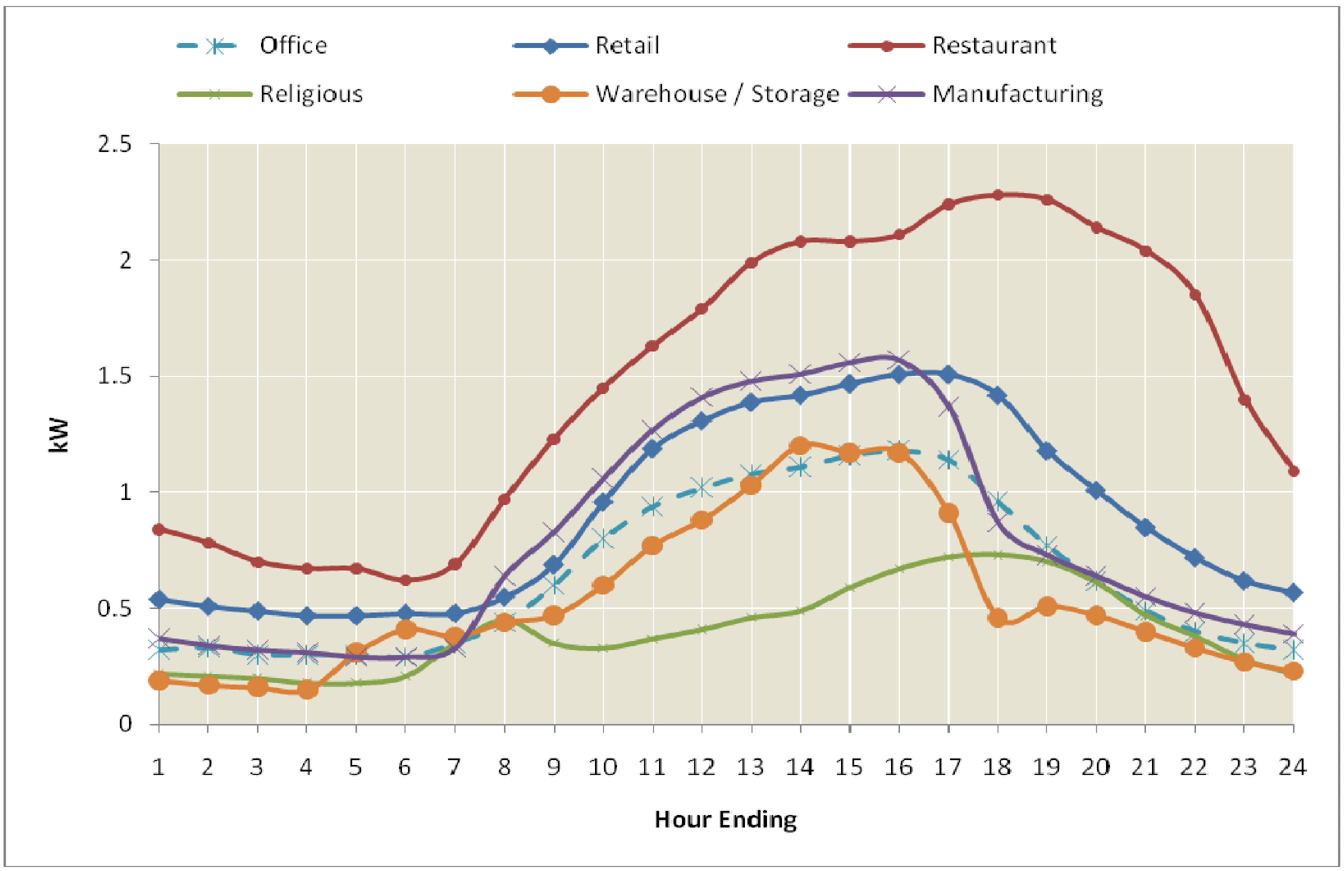

Figure 5. Average Hourly per AC unit Loads By Industry (I)

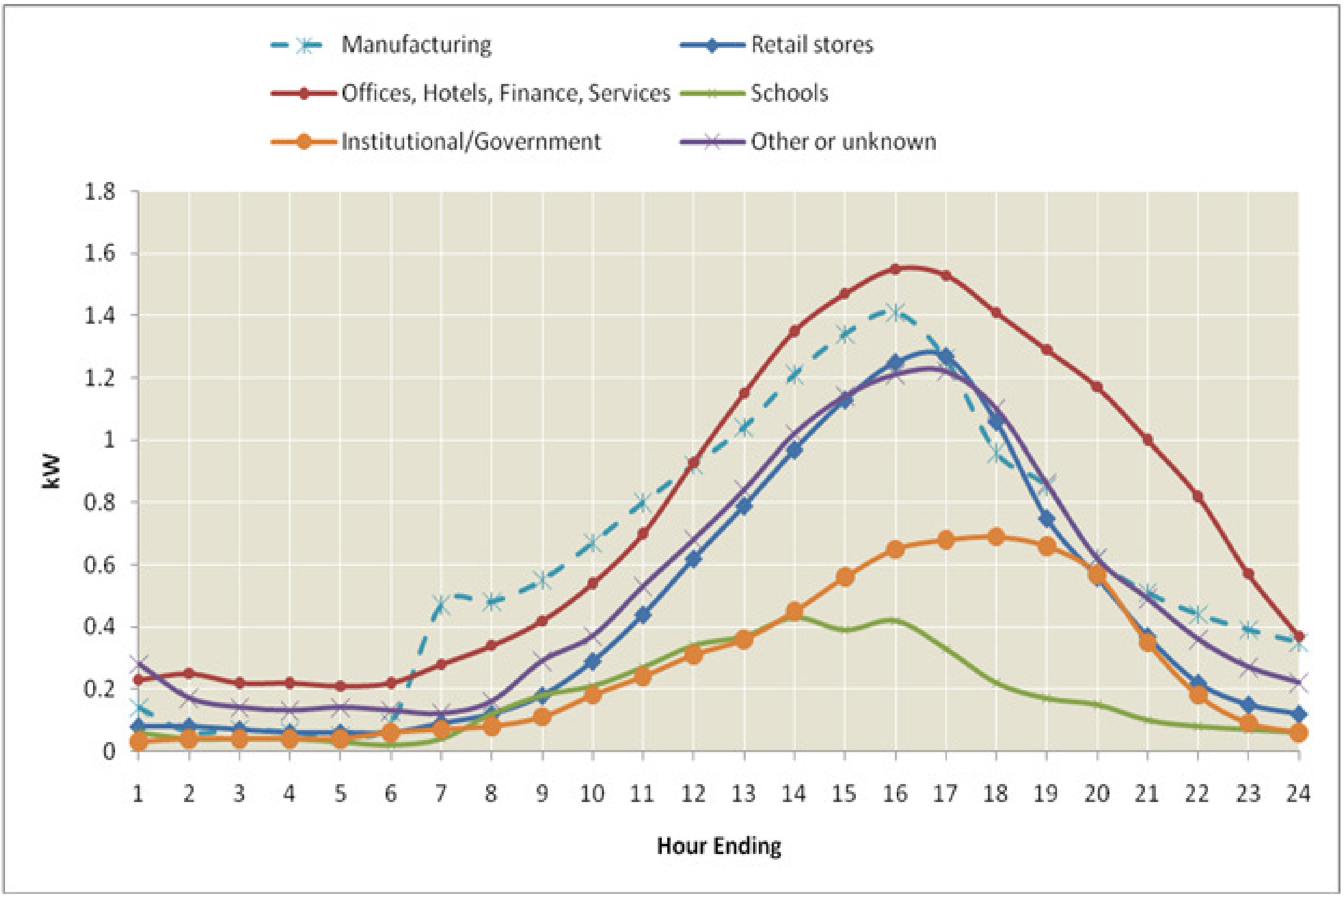

Figure 6. Average Hourly per AC unit Loads By Industry (II)

Conclusion

Commercial customers are potentially very attractive for AC load control programs due to their substantial AC use even at moderate temperatures, due to the fact that their peak loads occur earlier in the day and due to potential efficiencies in recruiting them. There are several issues that arise in cost-effectively recruiting commercial customers and measuring their load impacts. For recruiting, it is important to have correct contact information, to be able to reach the correct decision-maker at the business and to target only customers who are actually eligible and likely to provide adequate load impacts. In measuring load impacts, there is always a trade-off between cost and precision, and for commercial customers the cost of a given level of precision will be higher than for residential customers.

References

George, Stephen and Mike Perry, 2010. "2010 Residential and Small Commercial Demand Response Initiative and Hydro Ottawa *peaksaver*® Small Commercial Pilot Program Evaluation." FSC Group.

George, Stephen, Mike Perry and Peter Malaspina, 2010. "2010 Load Impact Evaluation for Pacific Gas and Electric Company's Smart AC Program." FSC Group.

George, Stephen, Mike Perry and Sarah Woehleke, 2010. "2010 Load Impact Evaluation of San Diego Gas & Electric Company's Summer Saver Program." FSC Group.

George, Stephen, Josh Bode, Mike Perry, Josh Schellenberg, and Sam Holmberg, 2010. “Load Impact Estimates for SCE's Demand Response Programs: Residential and Commercial Summer Discount Plan; Agricultural and Pumping Interruptible Program; Real Time Pricing.” FSC Group.