Results from California's Small Commercial Demand-Responsiveness Pilot Program

Roger L. Wright, RLW Analytics, Inc., Sonoma, CA
Mark S. Martinez, Southern California Edison, Rosemead, CA

ABSTRACT

Southern California Edison (SCE) initiated the SCE Energy$mart Thermostat^{SM} Pilot Program under a March 2001 order of the California Public Utilities Commission. The program is designed to test demand responsiveness and provides small commercial customers in SCE's service territory with two-way communicating thermostats. SCE curtails the AC load of the participants during critical periods by sending out a radio signal that instructs the thermostat to raise the cooling set point by a specified number of degrees, thereby reducing the cooling load until the new set point is reached.

This paper describes the results of the impact evaluation of the 2002 program conducted by RLW Analytics (RLW). The evaluation was informed by program tracking data, five-minute AC logger data for a small number of sites, and hourly run-time data collected from each thermostat. The logger data was used to estimate the hourly kW load of each AC from its hourly run-times, to validate the results, and to provide an independent estimate of curtailment impacts.

Nearly 4,500 thermostats were installed in 2002, the first operational year of the program. Small, one-stage ACs dominated the program, with 88% of the total tonnage. The typical four-degree, two-hour curtailment yielded a maximum reduction of approximately 10 MW, first hour energy savings between 6 and 7 MWh, and second hour energy savings of about 3 MWh. The average effective duration of the savings was about 55 minutes, and little or no snapback was seen following the curtailment period. The most important finding for future evaluations was that, given good information about unit tonnage, the thermostat run-time data provide accurate information about the hourly kWh loads and load reductions for the one-stage ACs that dominate the program.

Program Overview

The CPUC directed SCE to implement the SCE Energy$mart Thermostat^{SM} Pilot Program to test the viability of demand responsiveness among small commercial customers through two-way communicating thermostats. The program provides small commercial customers in SCE's service territory with two-way communicating thermostats. SCE uses a radio signal to remotely curtail the AC load of the participants during critical periods. When the curtailment is activated, the thermostat raises the cooling set point by a specified number of degrees called the temperature offset, thereby reducing the cooling load. The thermostat sends a radio signal back indicating it has received the signal and has implemented the temperature rollback. The thermostat reports any overrides by the participants and can collect and report the hourly run time of the controlled units.

Each participant received one or more free thermostats (including installation) and a $300 annual incentive per thermostat for participating in the pilot program. The participant is penalized $5 if the participant chooses to override a particular curtailment. A maximum of fifty curtailments could be called from May 1 through October 31 between noon and 6pm excluding weekends and holidays. Each curtailment could last no longer than four hours and have an offset no larger than 4º F.

In 2002, fourteen curtailments were called at a variety of times and weather conditions. SCE attempted to call curtailments in the middle of hot periods based on five-day weather forecasts, in order to simulate the conditions that might cause an ISO event. SCE also tried to match the timing of the curtailments to the ISO peak load. The SCE project manager typically initiated a curtailment on days when the predicted high temperature was greater than 90°F.

Program Objectives

The objectives of the program were specified by the CPUC in Decision 01.03.073 dated March 27, 2001. The program was to equip commercial buildings with the capacity to automatically control thermostats, evaluate installation of local infrastructure, and provide incentives for demand response. In addition, the CPUC was interested in consumer participation and behavior patterns in the program, consumer satisfaction with newer interactive demand response technologies, responsiveness of small commercial customer load to price or system demand signals, and the ability of such programs to deliver reliable and verifiable energy and demand savings. The CPUC provided SCE with general instructions on program administration, the technology to be used, the types of marketing they should utilize, and the maximum budgets. The CPUC instructed SCE to outsource implementation and evaluation and to implement the program as rapidly as possible, with a goal of approximately 5,000 small commercial customers and an estimated 4 MW of load under control before the end of 2002.

Eligibility

This pilot program targeted small commercial and nonprofit customers with less than 200 kW maximum demand and 1,000,000 kWh annual energy usage. The program excluded residential customers as well as other non-commercial accounts such as government, military, agricultural, and schools. SCE also excluded national or regional chain accounts, franchises, time-of-use accounts, other special rate accounts, participants in existing load management programs, and locations without coverage by the radio communication system used in the program.

Evaluation Overview

RLW contracted with SCE to conduct an impact evaluation of the 2002 program. The primary goal of the impact evaluation was to estimate the maximum kW load reduction, the effective duration of the load reduction, the kWh reduction during each hour of the curtailment period, and the snapback – the change in kWh consumption during the first hour following the curtailment period. Other objectives were to gauge the responsiveness of small commercial customer load to system demand signals, quantify the effect of time and temperature on demand reduction, and determine the reliability of data sources for planning future evaluations.

Data Sources

In addition to implementing the program itself, SCE collected extensive data during the summer and fall of 2002:

- Program Tracking Data, describing the thermostats and controlled air conditioners of all 4,325 program participants,
- Curtailment event summaries, characterizing the date, time and temperature offset of each curtailment during 2002,

- Information about signal reception and customer overrides collected from all installed thermostats,
- Fifteen-minute measurements of exterior temperature throughout the summer of 2002 collected from a central weather station representative of the participants,
- Thermostat Run-Time Data, collected on ten days from virtually all program participants,
- Five-minute AC logger data collected at early installations at sixteen offices at the Brackett Business Park (BBP) and 23 small restaurants called the Mini-pilot program (MPS).

Program Tracking Data. The program tracking system provides data for each thermostat installed in the program. At the time of this study, the SCE database described 4,325 installed units. The tonnage of the controlled air conditioning was provided for 2,566 of these units. We estimated that the 4,325 installed thermostats controlled a total of 17,994 tons of air conditioning.

The controlled ACs varied in size from 2 to 24 tons. The MPS logger data showed clearly that the larger units, i.e. 7 tons or larger, were all multi-stage units. These units were of concern because the thermostat run-time data were not designed to describe the operation of multi-stage units. However the smaller single-stage units represented 95% of the units and 88% of the total tonnage.

Thermostat Run-Time Data. A key element of our analysis strategy was the run-time data that can be retrieved from each of the Carrier EMi thermostats installed under the program. The system works as follows:

- Each hour, the thermostat monitors the equipment run time cycles and minutes as well as the average room temperature and set points.
- The thermostat can store the data for 24 hours for 7 days.
- Each thermostat can transmit the stored data on request to a central server.
- The central server stores the data and can serve it up for analysis.

Carrier claims that the run-time data provide an effective alternative to end use metering the AC load. If the run-time data are coupled with estimates of the operating kW of each AC unit, these data can be used to estimate the hourly kWh load of each installed unit, and thereby the hourly load reduction during each curtailment.

AC Logger Data. SCE and its contractors collected five-minute logger data from June through September measuring the instantaneous kW load of 39 controlled air conditioner units installed early in the program. SCE and its subcontractors collected the five-minute logger data for each of the following two experimental groups:

- Brackett Business Park (BBP) - sixteen thermostats and controlled AC units at thirteen small offices at a single office complex.
- Mini-Pilot Program (MPP) - 23 thermostats and controlled AC units at nine small restaurants.

We used data visualization software (Visualize-IT) to review the five-minute logger data for each of these 39 units. We made extensive use of these data in the statistical analysis.

Statistical Analysis Methodology

We used the logger data together with the fifteen-minute measurements of exterior temperature to estimate the load impact of curtailments called on nine different days during the summer of 2002. We treated the 39 AC units with logger data as if they were a sample of all installed AC units. Although these units were certainly not randomly selected from the full population of program participants, we post-stratified the 39 units according to the size of each AC unit, using the available tracking information about the size, in tons, of the population of all installed units. Using the case weights derived from the post-stratification, we estimated the average five-minute load per unit in the population of all participating AC units throughout the summer of 2002.

The impact of each of the nine curtailment events was developed by comparing the average five-minute kW load per unit on the curtailment day to the load on one or more non-curtailment, baseline days. The baseline days were selected to have comparable temperature conditions to the curtailment days. When necessary, the baseline load profile was adjusted to have comparable kWh in the two hours preceding the curtailment event. Using the difference between the load profiles on the curtailment day and baseline days, we calculated the maximum kW load reduction and the effective duration of the load reduction - defined to be the number of minutes during the curtailment period in which the load reduction was greater than one-half of the maximum load reduction. We also calculated the kWh reduction during each hour of the curtailment period and the snapback - the change in kWh during the first hour following the curtailment period.

We also made extensive use of the thermostat run time data. SCE and its subcontractors collected hourly run-time data for almost all of the participating units for a set of ten days in August through November. Our first step was to validate the accuracy of the hourly thermostat run times and to develop a procedure for converting the data into hourly kW load. We used the five-minute logger data for each of the sixteen units in BBP for this work. We started by calculating hourly run times from the five-minute logger data and comparing them to the run times reported by the thermostats. Then we used the five-minute logger data for each of the sixteen units to develop a statistical regression model that estimated the operating kW load of each unit as a function of its size in tons and the exterior temperature. We used this model to convert the hourly run times into estimates of the hourly kWh load of each of these sixteen units and compared the results to the hourly kWh load of each unit calculated from the five-minute logger data.

Finally, we used the run-time data collected for the full population of participating units together with the tracking information about the size of each unit to estimate the average hourly kWh load of all participating units on the ten days in which the run-time data had been collected. The ten days included three curtailment days, six days that were considered to be comparable non-curtailment days and one day in November that was excluded from the analysis. We used these data to provide an independent estimate of the hourly kWh impact among all program participants of the curtailments on these three curtailment days. Last, we extrapolated the results by calculating the impact per ton and comparing the results to the estimated total tonnage of all installed units.

Estimating Impact using the Logger Data Directly

In this section we discuss our analysis of the average load from the 39 logger units to estimate the impact of curtailments. We will start by using the August 27 curtailment as an example. The August 27 curtailment was a four-degree setback from 3 to 5 PM.

Our first step was to choose a comparison day. Figure 1 shows the temperatures for August 27 and ten nearby days. August 27 was a Tuesday and the maximum temperature was a moderate 86.4 degrees. August 23 appeared to be a good comparison day except that it followed a relatively cool day. August 26 could be considered as a comparison day but had a curtailment so it is not appropriate. August 28 appeared to be the best comparison day.

Figure 2 shows the actual load on August 27 and the baseline load, i.e., the load that might have been expected in the absence of a curtailment. To develop the baseline load, we adjusted the actual load on August 28 to better reflect the load of August 27 by multiplying the load on August 28 by a fixed factor called the true up adjustment. We calculated the true up adjustment as the ratio between the average load during the two hours prior to the curtailment on August 27 divided by the average load during the corresponding two hours on August 28.

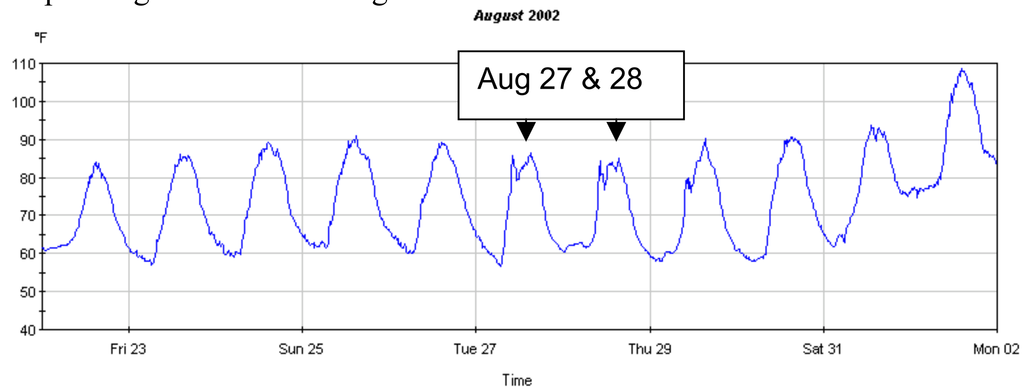

Figure 1: Picking Potential Comparison Days for August 27

Figure 2 also shows the estimated impact, i.e., the difference between the actual and baseline load. The graph shows that initial impact was about 2 kW per air conditioning unit, but the impact diminished to nearly zero by the end of the two-hour period.

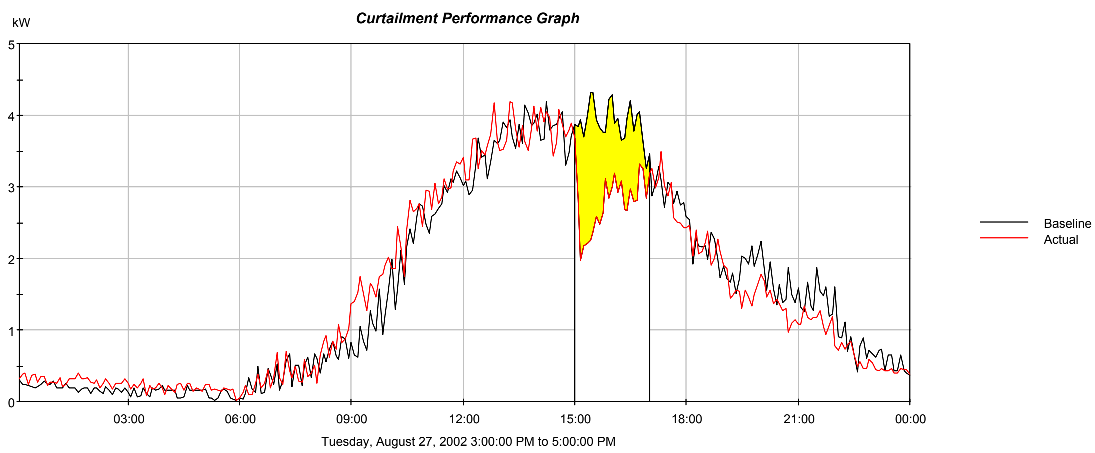

Figure 2: Estimating the Impact on August 27

Figure 3 gives the summary characteristics of the event. The summary shows the date and time of the curtailment and the temperature offset. The table also shows the high temperature on the day of the curtailment, as well as the override rate. The override rate is the number of the units that overrode the curtailment as a percentage of the units that confirmed the call. This statistic is taken from the event tracking data maintained by the Carrier system for the BBP and MPS control groups. In this case the override rate was 22%. The remaining statistics in Figure 3 reflect our analysis of the average load of the 39 units with logger data. The summary shows that the true up adjustment was about 111%. In other words, the load on August 28 was increased by 11% to provide a more accurate representation of the load on August 27 during the two hours prior to the curtailment. During the curtailment, the average load per unit dropped by a maximum of 2.08 kW. The minimum reduction was 0.27 kW during the period, indicating that the savings did not persist through the full period. The energy savings during the entire curtailment period was estimated to be 2.28 kWh.

We estimated the effective duration of the savings using a 'half life' concept. We defined the effective duration to be the number of minutes during the period that the kW reduction was greater than one half of the maximum reduction. Using this measure, we calculated that the reduction had an effective duration of 65 minutes. We also calculated the kWh savings during each hour of the curtailment, as well as during the first hour following the curtailment period. From the first hour to the second hour of the curtailment period, the energy savings fell from 1.5 kWh to 0.8 kWh. This is consistent with the estimated effective duration of the curtailment. In the hour following the curtailment period, the energy decreased by 0.1 kWh. The small decrease in energy in the hour following the curtailment suggests that there was little or no significant snapback.

Curtailment Date	8/27/2002
Start Time	3:00 PM
End Time	5:00 PM
Offset	4 Degrees
Curtailment Day High	86.4
Override Rate	22%
Baseline Day(s)	8/28/2002
Baseline Day High	85.1
Trueup Adjustment	111.4%
Maximum Reduction	2.08kW
Minimum Reduction	0.27kW
Energy Savings	2.28kWh
Minutes Duration	65
kWh Savings in Hour 1	1.5
kWh Savings in Hour 2	0.8
Hour Following	0.1

Figure 3: Summary Report for August 27

Summary of "Four-Degree" Curtailment Days with Logger Data

The analysis described above was carried out for a total of nine curtailment days of which seven were four-degree curtailments. The two remaining days were excluded from the summary since they were not considered to be comparable to the primary seven days. On one of these two days, a 2-degree curtailment was called. On the other day, a 2-degree curtailment was initially called but then changed to a 4-degree curtailment when a stage-two alert was announced.

Based on a simple numerical average, the typical four-degree event can be described as follows:

- The maximum reduction was 2.2 kW,
- The savings in the first hour was 1.4 kWh,
- The savings in the second hour was 0.7 kWh,
- The effective duration was 55 minutes, measured as the 'half-life' of the savings,
- The savings was generally very small after two hours,
- There was little or no snapback following the curtailment period.

Using the case weights previously discussed, we found that the average size of the sample units was 3.93 tons. We used this to convert the impact from kW per unit to kW per ton as shown in Figure 4. We also estimated the total program-wide impact of a four-degree curtailment. As reported earlier, we estimated that the 4,325 participating units have a total size of 17,994 tons. Using this information together with the kW impacts per ton found in the sample, we estimated the total program impact as shown in Figure 4.

From this analysis, we estimated that among all 4,325 participating units, a four-degree, two-hour curtailment will yield a maximum kW reduction of approximately 10 MW, first hour energy savings of about 6 MWh, second hour energy savings of about 3 MWh, and negligible savings at the end of two hours.

	Per Sample Unit	*Per Sample Ton*	*Program Total*
Maximum kW Reduction	2.23	0.567	10,199
kWh Savings in Hour 1	1.4	0.354	6,362
kWh Savings in Hour 2	0.7	0.187	3,371

Figure 4: Estimated Program-Wide Impact Using the four-degree Logger Results

Analysis of the Thermostat Run-Time Data

During 2002, SCE collected hourly thermostat run-time data from all installed thermostats throughout the program on several days in August through November. During the summer of 2002, SCE collected run-time data for the following ten days: August 26, **27,** 28; Sept 9, **10,** 11, 23: Oct 7 **8;** and Nov 20. The days in **bold** were curtailment days. The remaining days can be considered as potential baseline days.

The great advantage of the run-time data is that it is available from all responding thermostats. Moreover, the run-time data can be collected at relatively low cost because the thermostats themselves generate these data. The disadvantage of the run-time data is that the thermostats do not provide kW itself, only the number of minutes of air conditioner operation in each hour. The run time has to be converted to the kW consumption using information about the kW load of the air conditioners when they are operating. Moreover, the hourly data do not provide the fine resolution of the five-minute logger data.

Examining the Logger Data

In this analysis, we worked with the sixteen AC units with logger data collected in the BBP study. All of these units were one-stage ACs between two and five tons. Our analysis was designed to determine the best way to analyze the remaining units in the program that have run-time. We started by examining the five-minute logger data for each individual BBP site. We observed the following three especially significant characteristics.

Spikes. Spikes are unusually high kW measurements that last only momentarily. We observed a few units where the operating load was typically between 2 kW and 2.5 kW, but with a few isolated measurements that exceeded 3.5 kW – a full 1 kW higher than the other measurements. Because of their short duration, we felt that these spikes were spurious and should be deleted from the analysis.

Heating Load. Heat pumps can go into heating mode on some cold days. We noticed an unusually large load appearing in October for a couple of the BBP units. The peak load in October for one of those units was 3.5 kW whereas the typical peak load of this unit throughout the summer was in the range 2 to 2.5 kW. On October 10 and 11, the load was especially high during night as might be expected of a heating load. We concluded that this load was not characteristic of air conditioning operation and should be deleted from the analysis. We found other units that displayed this type of characteristic late in October. To avoid these loads, we limited our run-time analysis to June through September.

Operating load. The operating load is the kW draw of the unit when it is operating in cooling mode. The logger data showed that the operating load of these units varied, as expected, with their size as measured in tons. Somewhat more surprising to some, the logger data also showed that the exterior temperature significantly affected the operating load of each individual unit.

Validating the Run Times Reported by the Thermostats

Our next task was to use the logger data to calculate the run time of the unit, i.e., the number of minutes of operation during each hour. We started by looking for a suitable calculation method. For the typical two-ton AC, the measured values were either equal to zero, signifying that the unit was off, or roughly 2 kW when the unit was operating. For these units, the minutes of run time in a given hour should be approximately equal to five times the number of five-minute logger readings that are greater than zero during the hour.

However for several units, the logger data was either about 2 kW or about 0.25 kW, but was never equal to zero. These units appeared to have a fan that was programmed to run continuously during all weekday hours as well as during most of the weekend hours. The logger data showed that the fan schedules were not affected by the run time of the air conditioner. Moreover the fans were not affected by the curtailments. To ensure that these fan loads were not included in our calculated run time, we redefined our method for estimating the minutes of run time in a given hour to be five times the number of measurements that are greater than 1 kW during the hour.

Using this methodology, we used the logger data to calculate the run time for each individual unit in each hour and then compared the results to the run times reported from the thermostats themselves. For all sixteen units, the discrepancy was 3% or smaller and the correlation was 0.979 or higher. The average absolute value of the discrepancy was only 1% and the average correlation was 0.993. From these results we concluded that the run times reported by the thermostats are very accurate.

Estimating the Operating Load

In the preceding task we demonstrated that we could predict the run time reported by the thermostats from the hourly kW consumption recorded by the loggers. The next task was to address the reverse problem –to estimate the hourly kW consumption of each AC unit from the run time reported by the corresponding thermostat. The operating load of each unit was the key to converting the thermostat run-time data into an estimate of the kW load of each unit. We used the five-minute logger data to calculate the operating load of each of the sixteen BBP units in each hour from 2 pm through 5 pm in the months June through August. We used regression analysis to model the operating load as a function of the tons of the unit and the exterior temperature of the hour. In particular, we used the logger data to estimate a regression model relating the operating load per ton of each unit to the exterior temperature.

We started by examining the data graphically. We prepared scatter plots that indicated a strong linear relationship between the operating load of a unit and the exterior temperature. In some cases the graphs revealed spikes in the observed load. We screened the spikes from the analysis database by saving the residuals from a separate linear regression relating operating load to exterior temperature for each logger, and deleting the observation if the absolute value of the residual was greater than 0.1 times the tons of the unit. Using this filter, we dropped 28 out of 2,437 observations.

The next step of our analysis was to again fit a separate linear regression relating operating load to exterior temperature for each unit, this time using the screened data. We saved both the intercept and slope from the regression model associated with each individual unit. Then we regressed both the intercepts and the slopes on the size of each air conditioner, measured in tons. After some exploration we arrived at the simple predictive equations: Intercept = 0.5311 * tons, Coefficient = 0.0086 * tons.

These equations imply that the operating load per ton can be estimated as the following linear function of exterior temperature: 0.5311 + 0.0086 * (Exterior Temperature in F).

Figure 5 shows our resulting estimates of the operating load of a single stage AC unit versus the exterior temperature. As an example, a two-ton unit would be expected to have an operating load of about 2.62 kW on a 90-degree F afternoon.

Exterior Temperature	***kW / Ton***
70	1.13
75	1.18
80	1.22
85	1.26
90	1.31
95	1.35
100	1.39

Figure 5: Operating Load per Ton vs. Exterior Temperature in Degrees Fahrenheit

Estimating Hourly kW Load from Thermostat Run-Time Data

This task completed our study of the validity of the thermostat run-time data. In this task we estimated the hourly kW load of each of the sixteen BBP units from the thermostat run time data and then compared the results to the hourly kW load calculated directly from the five-minute logger data. We used the model developed in the prior section to estimate the operating load in each hour for each unit as a function of the minutes of run time in the corresponding hour as reported by the thermostat, the tons of the unit, and the exterior temperature during the hour.

We also calculated the actual kW in each hour for each unit as the mean of the measured kW in the logger data for each five-minute period in the hour. Finally we calculated the simple unweighted average hourly kW for all sixteen loggers using both the estimate and actual hourly loads of each logger. Using these data we found that the kW estimates were highly correlated to the actual kW measurements from hour to hour for all units. In particular, we found that the average hourly load using the thermostat information was equal to 1.27 kW across all hours, which was only 0.05 kW, or 4%, larger than the observed load. Moreover the correlation was extremely high, 0.997. As an example of the accuracy of the thermostat run-time data,

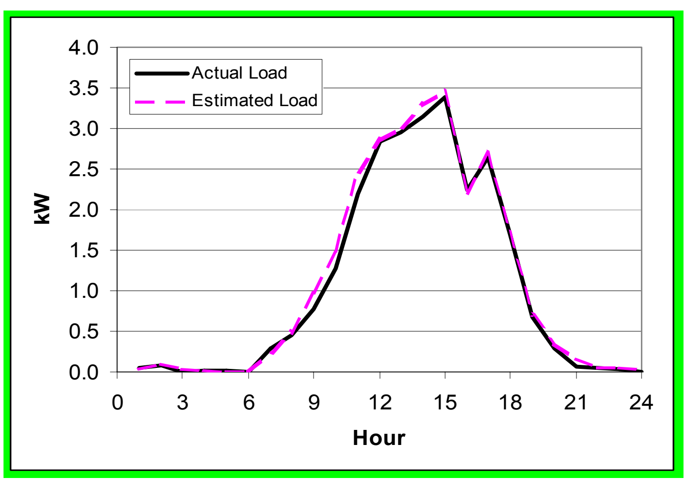

Figure 6: Actual vs. Estimate Average Hourly Load of all sixteen BBP Units on August 27

Figure 6 compares the estimated and actual average hourly load of all sixteen BBP units on August 27. These results demonstrate that the run-time data provided by the thermostats yields a very accurate estimate of the average hourly load, at least in a relatively small set of smaller, one-stage units.

Estimating Impacts using the Run-Time Data

The great strength of the thermostat run-time data is that it measures the response of almost all of the thermostats installed in the program. Therefore the results are likely to be representative of all participating units. Moreover, we have demonstrated that the run-time data can be used to develop an accurate estimate of the average load of single-stage units. In addition, these data are inexpensive compared to interval metering. In this section we will apply the methods discussed in the preceding section to analyze all of the run-time data available for the summer of 2002.

Recall that there were a total of 2,566 participating units with recorded tons in the tracking database and that these units have an average size of 4.16 tons per unit. Our analysis was limited to the units with reported run time data as well as known size. We found that that 2,187 units, i.e., over 85%, had reported run time data as well as recorded size. The 15% with missing run time data was due to communication or other problems. These 2,187 units with both recorded tons and run time data had an average size of 4.15 tons per unit, practically identical to the average size of the 2,566 units with reported size. As shown in Figure 7, the smaller single-stage units comprise 95% of all units and 88% of the total tons.

Size in Tons	*Number*	*Percent*	*Total Tons*	*Percent*	*Tons per Unit*
6 or Smaller	2,076	95%	8,021	88%	3.864
7 or Larger	111	5%	1,061	12%	9.559
Total	2,187	100%	9,082	100%	4.153

Figure 7: Summary Statistics for the Units Included in the Run Time Analysis

We also compared the units included in the run-time analysis with the units included in our prior analysis of the five-minute logger data. In the prior analysis, we worked with the 39 units from the BBP and MPS studies. The average size of those 39 units was 3.93 tons per unit. By contrast, the average size of the units in the run-time analysis was 4.15 tons per unit. Therefore these units are, on average, about 6% larger than the units included in the logger analysis.

We can also compare the run-time results for August 27 with the logger results reported earlier. Recall that the curtailment on August 27 was a four-degree offset from 3 PM to 5 PM. As in our previous analysis of this event using the logger data, we have selected August 28 as the comparison day. Figure 8 displays the results of our new analysis using all available run-time data for installed units. This graph can be compared to Figure 2. Note that the resolution in Figure 8is cruder than in Figure 2 since the run time data is hourly whereas the logger data was on a five-minute interval.

Figure 9 summarizes the results of the new analysis. This table is similar to Figure 3 developed from the five-minute logger data, but some of the statistics such as minutes duration are no longer reported since the current results are based on hourly load. The statistic labeled Confirmed Units reflects the number of units that were operable on August 27 and confirmed the curtailment signal in the entire system-wide group. The override rate is based on all confirmed operable units in the system-wide group. The table also shows the number of units included in the run time analysis of the August 27 event, in total and as a proportion of all confirmed units in the system-wide control group. The remaining characteristics reported in Figure 9 are equivalent to their use in the summary reports that were presented in Figure 3.

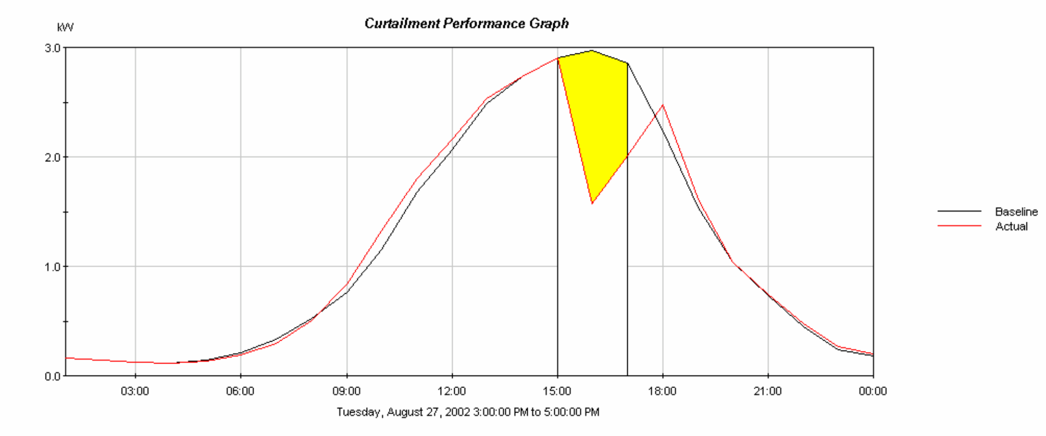

Figure 8: Estimating the Impact on August 27 using the Run-Time Data

Curtailment Date	8/27/2002
Start Time	3:00 PM
End Time	5:00 PM
Offset	4 Degrees
Curtailment Day High	86.4
Override Rate	16%
Confirmed Units	2,569
Units in the Analysis	1,872
Proportion of Confirmed	73%
Baseline Day(s)	8/28/2002
Baseline Day High	85.1
Trueup Adjustment	107.0%
kWh Savings in Hour 1	1.4
kWh Savings in Hour 2	0.8
Hour Following	-0.2

Figure 9: Summary Report for August 27 from the Run Time Analysis

In comparing Figure 9 to Figure 3, we see that Figure 9 shows that the run time analysis is based on 1,872 units, or 73% of all 2,569 confirmed units. By contrast, Figure 3 is based on the 39 units in the BBP and MPS with five-minute logger data. The override rate within this group was only 16% whereas it was 22% within the BBP and MPS control groups. The savings in the first and second hours of the curtailment reported in Figure 9 are almost equal to the comparable statistics reported in Figure 3. However, it should be recalled that the units included in Figure 9 were about 6% larger than the units included in the logger analysis. Therefore, on a per ton basis, the new results are about 6% smaller than the prior results. Figure 9 suggests that there might be some small snapback in the hour following the curtailment whereas Figure 3 indicated no snapback. In short, the results based on the run-time data are remarkably similar to the results from the logger data for August 27.

During the summer of 2002, there were two curtailment days that had both logger data and run time data. Both days had four-degree offsets. Figure 10 compares the results of the logger and run time analysis for these two days. The data for August 27 are taken directly from Figure 3 and Figure 9, but we have shown the results to one more decimal place to facilitate the comparison. The data for 10/8/2002 have not previously been presented in this paper. The table shows that the two methods of analysis yielded remarkably similar estimates of savings in the first and second hours of these two curtailment events.

Curtailment Date	8/27/2002		10/8/2002	
Method	RTD	Loggers	RTD	Loggers
kWh Savings in Hour 1	1.40	1.45	1.54	1.57
kWh Savings in Hour 2	0.84	0.82	0.83	0.62
Hour Following	-0.23	0.06	-0.37	-0.19

Figure 10: Summary of Savings per Unit for four-degree Curtailments

Figure 11 provides a similar comparison on a per-ton basis. The top portion of the table shows the average size of the units included in the two types of analysis, measured in tons per unit. The lower portion of the table restates the results from Figure 10 in kWh savings per ton.

Curtailment Date	8/27/2002		10/8/2002	
Method	RTD	Loggers	RTD	Loggers
Average Size in Sample	4.153	3.932	4.153	3.932
kWh Savings per Ton in Hour 1	0.337	0.370	0.371	0.400
kWh Savings per Ton in Hour 2	0.202	0.209	0.200	0.159
Hour Following	-0.055	0.014	-0.089	-0.047

Figure 11: Summary of Savings per Ton for four-degree Curtailments

Figure 12 summarizes the characteristics of all units installed in the program. We can estimate the total impact of the program by multiplying the per ton impact shown in Figure 11 by the total tons of the installed units, 17,994. Figure 12 shows the results.

Curtailment Date	8/27/2002		10/8/2002	
Method	RTD	Loggers	RTD	Loggers
Total Tons in Program	17,994	17,994	17,994	17,994
MWh Savings in Hour 1	6.1	6.7	6.7	7.2
MWh Savings in Hour 2	3.6	3.8	3.6	2.9
Hour Following	-1.0	0.3	-1.6	-0.9

Figure 12: Estimated Total Impact of Program-Wide four-degree Curtailment

We conclude with the following observations. The new results shown in Figure 12 are consistent with the total program impacts developed earlier in our analysis of the logger data. Moreover the results for these two days are consistent with the logger results for the seven curtailment days summarized in Figure 4. From these results we conclude that:

1. The SCE Energy$mart Thermostat^{SM} program has surpassed the goals given to SCE for the program (approximately 10 MW maximum load reduction vs. 4 MW CPUC goal for SCE),

2. The run-time data collected from the thermostats can play an important role in future impact evaluations of this type of program. Augmented by tracking information about the tons of each AC and logger information about the operating load per ton, these data can provide reliable estimates of hourly kW load and the impacts of curtailments.

3. Five-minute logger data collected on a sample of units can be extremely valuable for estimating the operating load and providing higher-resolution information about load shapes and the duration of the impacts.