



Final Report Pacific Gas and Electric SmartAC Load Impact Evaluation

**Pacific Gas and Electric Company
San Francisco, California
April 24, 2008**

Table of Contents

- 1. Executive Summary 1-1
- 2. Introduction 2-1
 - 2.1 Scope of the Evaluation 2-1
 - 2.2 Program Description 2-1
 - 2.3 Report Organization 2-2
- 3. Program Details 3-3
 - 3.1.1 Technology Options 3-3
 - 3.1.2 Control Mechanism 3-3
 - 3.1.3 Thermostats 3-6
- 4. Data and Methodology 4-9
 - 4.1 Data 4-9
 - 4.1.1 Program Enrollment Data 4-9
 - 4.1.2 Load Data 4-10
 - 4.1.3 Weather Data 4-16
 - 4.2 Overview of Analysis Methodologies 4-18
 - 4.2.1 Load Model 4-18
 - 4.2.2 Duty Cycle Modeling 4-20
 - 4.2.3 Choice of Weather Data 4-23
 - 4.2.4 Impact Estimates 4-23
 - 4.2.5 Projected Impact Estimates for General Conditions 4-24
 - 4.2.6 Ratio Estimation 4-25
 - 4.3 Drivers of the Impact Results 4-27
 - 4.3.1 Participation-Related Factors 4-27
- 5. Impact Evaluation Results 5-29
 - 5.1 Event Level Results 5-30
 - 5.2 Impact at Time of System Peak 5-31
 - 5.3 Comparison of Thermostat Ramping Strategies 5-35
 - 5.4 Load Reduction by Site and Unit Characteristics 5-37
 - 5.4.1 Large vs. Small AC Units 5-37
 - 5.4.2 Single vs. Multiple AC Units 5-39
 - 5.5 Projections for 2008 5-41
 - 5.5.1 Switch Projections for 2008 5-41
 - 5.5.2 Thermostat Impact Projections for 2008 5-44
 - 5.6 Connected Load 5-47
 - 5.7 Snapback 5-47
- 6. Conclusions 6-1
 - 6.1.1 Control Strategy 6-1
 - 6.2 Participant Satisfaction 6-1
 - 6.3 Switch Summary 6-2
 - 6.4 Thermostat Summary 6-2

Table of Contents

7.	Appendix A	7-1
7.1	Tobit Model of Duty Cycles	7-1
7.2	Individual Demand Impacts by Duty Cycle Approach.....	7-3

1. Executive Summary

This report contains KEMA's load impact evaluation of PG&E's SmartAC Program after its first year of operation (2007). The results provided in this report show that the SmartAC program successfully reduced air conditioning load during the event periods.

Program Description

PG&E's SmartAC Program is a direct load control (DLC) program that first began enlisting customers in spring 2007. The Program began by recruiting customers in San Joaquin County (the city of Stockton and its surrounding areas) but has since expanded to other areas of the PG&E service territory. As of the end of the 2007 cooling season, the program had over ten thousand participants. The vast majority were residential customers.

The SmartAC Program uses paging signals to reduce the energy consumption of participants' air conditioners during times of peak system demand. The air conditioners are controlled either by a programmable thermostat or a switch that the Program installs at the participant's residence or business. The switch employs an adaptive switching technology that controls the air conditioner based on prior air conditioning behavior. Of the 8,800 participants enrolled as of the end of August, 2007 (at the time of system peak), 30 percent opted for the programmable thermostat while the remainder opted for the switch.

Load Impact Evaluation Methodology

The load impact evaluation's primary goal was to estimate residential demand reduction as a function of temperature, time of event, event duration, and unit size. KEMA selected a sample of 297 homes with 353 AC units for the metering sample. The sample was split into two roughly equal groups for the two control technologies, further stratifying by size (cooling tons from all units) and number of units (one, or two or more.) Models for AC unit-specific baselines were developed for load and duty cycle, and used to compare with event day performance. KEMA aggregated unit-level results using a ratio estimator on tons to improve accuracy and address potential sample bias.

This evaluation took place at a time of daily, and substantial, increases in the number of SmartAC participants. Participation increased from 6 to over 10 thousand residential customers during the cooling season of 2007. The weights utilized for the analysis varied accordingly - for each event, they were based on the composition of the population on that day.

The SmartAC load impact evaluation included fifteen events. Of those, two were conducted for the entire population of SmartAC participants, and the remainder were conducted for the sample only. The two population events occurred under relatively mild conditions.

Load Impact Evaluation Results

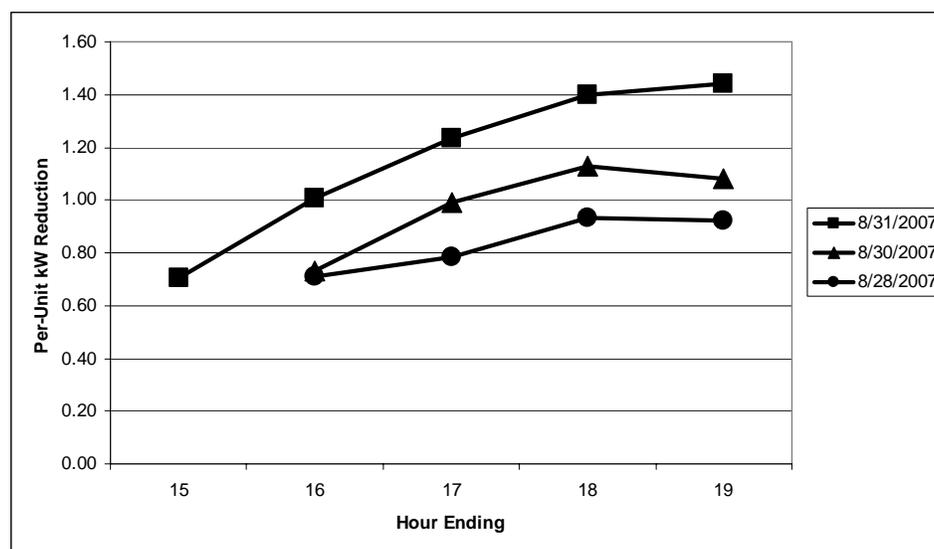
The SmartAC evaluation results indicate substantial demand reduction for the participants with switches. Participants with thermostat controls, facing the ramping strategies employed by the Program in 2007, produced demand reduction lower than that realized by switch participants. The results for the two

control technology groups reflect potential differences from self-selected populations¹ and *ex ante* control levels. Given the populations as sampled and the control levels employed, the switches provided greater demand reduction. These results should not be construed as necessarily reflecting the relative efficacy of thermostat controls versus switches.

Overall, Program 2007 savings are estimated to range between 0 and 1.21 kW per unit (average over the duration of the SmartAC event, depending on weather conditions at the time of the event.) Thermostat savings range between 0 and 0.89 kW per unit, and switch estimates between 0 and 1.34 kW per unit. A number of the events called for the purpose of this evaluation were on relatively mild days. On these days, effectively no savings were identified. These results include the effect of participants opting out of events. Opting out of an event was an option for SmartAC participants but was almost never used.

Figure 1-1 illustrates per-unit savings on the three hottest days of the load impact evaluation. On August 28, 2007, the average hourly temperature was 84, and the high temperature of the day was 99. These temperatures were 87 and 100, and 88 and 98, respectively, for August 30 and August 31. The maximum load impact estimated is 1.44 kW per unit, on August 31 at 7 PM.

Figure 1-1
Program Impact Results, Average Per-Unit kW Reduction per Hour on the
Three Hottest Days of the Load Impact Evaluation



The day of the system peak (August 29, 2007 at 5 PM) was not an event day. Using impact estimate from the surrounding days, full program savings at time of the system peak are estimated to be between 0.88 and 1.23 kW per unit. The low estimate of savings, 0.88 is the average 5 PM savings of the two surrounding days. The high estimate comes from the event day two days after the peak when the temperature was closest to the peak day’s conditions. Had a Program-wide event been activated on that

¹ Program participants at the time of this evaluation were exposed to information regarding both technologies, and thus self-selected into one of the control device populations. It is very likely that the choice of technology may be a proxy for behavioral differences that are not accounted for in this evaluation.

day, this would have translated into peak savings of 7.8 to 10.9 MW from the Program participants that were enrolled at the time.

Two different setpoint increase ramping strategies were compared. Participants with thermostat control were divided into two groups and received the different ramping strategies in alternating fashion. The steeper ramping strategy was expected to provide increased load reduction at least in some hours. Results confirm this expectation though apparent differences between the ramping strategy groups make this a tentative conclusion.

This evaluation also looked at impact per ton relative to unit size and number of units for the two different control technology groups. These results are suggestive but impossible to disentangle from other possible group differences.

Load Impact Projections

The load impact models were utilized to project savings across a range of temperatures.

Savings for SmartAC Switches are estimated to range between 0.3 and 0.7 kW per unit on days when the average temperature is 80 degrees, to 1.54 and 1.74 kW per unit on days when the average temperature reaches 94 degrees.

Savings for SmartAC Thermostats are estimated to range between 0.3 and 0.4 kW per unit on days when the average temperature is 80 degrees, to 1.3 and 0.5 kW per unit on days when the average temperature reaches 94 degrees. The Thermostat load impact estimates derived from this model are lower than the Thermostat load impact estimates derived from day of the event data. This suggests the need for model refinements that will be possible when additional data becomes available in 2008.

Program level projections are not provided, because they are determined by the number of each type of SmartAC device installed at the time of the event.

2. Introduction

2.1 Scope of the Evaluation

The impact evaluation goals for the SmartAC program include:

- Estimates of demand reduction as a function of temperature, time of event, event duration and unit size.
- The temperature sensitivity of demand reduction and connected load.
- Effects of customer behavior, override, signal/switch failure, attrition and snapback.
- Measuring the difference between two different thermostat setback approaches²
- Avoiding bias despite a likely non-representative sample³.

To meet these goals, three hundred residential customers were recruited for a measurement and verification sample, and fifteen load control events were called across a wide range of temperature conditions. These data provide insight into demand reduction as a function of temperature, time of event, event duration and unit size. We also use these data to provide projected demand reduction and connected load across a range of temperature conditions. All results are reported by the two control types and thermostat control results are also reported by ramping strategy. We used a ratio estimation approach to address concerns about bias due to the ongoing development of the program population.

There was not enough over-ride behavior or attrition to make possible any conclusions as to the potential effect of these important drivers of demand reduction. These issues will be revisited as the program matures.

2.2 Program Description

PG&E's SmartAC Program is a direct load control (DLC) program. The program uses paging signals and control technologies to limit air conditioner usage during program events. Actual program events are triggered by insufficient system operating reserve or localized transmission or distribution emergencies constraints. The events used for this evaluation were called for the purposes of the evaluation itself.

The SmartAC Program first began enlisting customers in Spring 2007. As of the end of August, 2007 the program had approximately 8,800 participants. By January, 2008 there were 26,000 participants with installed devices and another 22,000 participants who were enrolled and awaiting device installations. The vast majority of these participants are residential customers. The Program began by recruiting customers in San Joaquin County (the city of Stockton and its surrounding areas) but has since expanded to other areas of the PG&E service territory.

² This evaluation goal was added after project kick-off, on May 31, 2007.

³ The EM&V sample for the first year of the SmartAC program was recruited from the first three thousand participants. By the end of 2007, the program had expanded to twenty-five thousand participants. While the sample was representative of participants at the time it was recruited, it is very likely that the characteristics of participants changed as participation in the program increased.

PG&E manages the marketing efforts, initiates the control events, and manages the overall program. An implementation vendor handles the dedicated Program hotline, enrolls customers, schedules installation appointments, and installs the control devices. A technology vendor manufactures the control devices and manages the system that makes the control events possible. The Smart AC Program pays all participants a one-time \$25 “thank you” payment.

2.3 Report Organization

Section 3 provides an in depth discussion of the program details relevant to this impact evaluation. In particular, these include the choice of the two control technology options and the implications these technologies have for load impacts. Section 4 provides a discussion of the data and methodology used for this analysis. Section 4 includes a discussion of drivers of impact results that provides further background for the results section. A further discussion of methodology is included in an appendix. Section 5 provides the results from the analysis. Section 6 provides conclusions.

3. Program Details

PG&E's SmartAC program is a complex, large-scale endeavor. Many aspects of the program have already been discussed and evaluated by KEMA in an earlier, process-oriented report⁴. This impact evaluation is focused on the program performance in terms of load reduction. Aspects of the program design have a direct bearing on the potential impacts of the program. This section discusses the specifics of the SmartAC Program as they relate to these potential impacts.

3.1.1 Technology Options

The SmartAC Program offers participants a choice of control technologies for their AC units. The thermostats and switches used by the PG&E SmartAC program share only basic similarities. Both technologies use one-way communication capabilities to remotely control usage at the AC unit. Each technology can be communicated with via a paging device and activated and de-activated for a program event. Neither the switches nor the thermostats have two-way communication capabilities.

Beyond these basic similarities, the two control technologies offer a quite different array of characteristics. The thermostat option offered by the program replaces a participant's existing thermostat(s) with a programmable thermostat advertised as a \$200 value. The participant gains the functionality of a programmable thermostat, including web-based access to remotely change the home AC settings. Installation of the thermostat requires an indoor visit by a technician.

The switch technology offered by the SmartAC program is installed at the AC unit. The installation of the switch generally does not require an indoor visit. The switch provides the participants no additional functionality. The switch may be effectively invisible to the participant.

The different features of these two control technologies give PG&E flexibility in their marketing efforts. Participants may be motivated by monetary value and functionality (thermostat) or privacy and invisibility (switch). A choice of control technologies allows PG&E to cater to a wider range of potential participants.

3.1.2 Control Mechanism

Thermostats and switches control AC usage in different ways. In general, switches control AC unit compressor run-time while thermostats control household temperature.

3.1.2.1 Switches

Switches directly control operation of the AC compressor. Traditional switches control AC usage by limiting potential compressor run-time to a maximum amount during a time period. For a 50% control level, for instance, the unit will only be able to run a maximum of half of the time. Most commonly, programs with 50 percent control level limit run-time to 15 minutes per half hour (50% * 30 min = 15 minutes of control). All avoided run-time above 15 minutes per half hour represents load reduction. This could amount to 15 minutes of impact per half an hour if the unit would have been running full time.

⁴ Final Report: Process Evaluation of 2007 PG&E Smart AC Program. Study ID PGE0262.01. March 31, 2008

Load reduction will be less if the unit would have been running less than full time. A unit running at only 50 percent run-time, however, will adjust to the schedule enforced by the switch but will provide essentially no load reduction for the program.

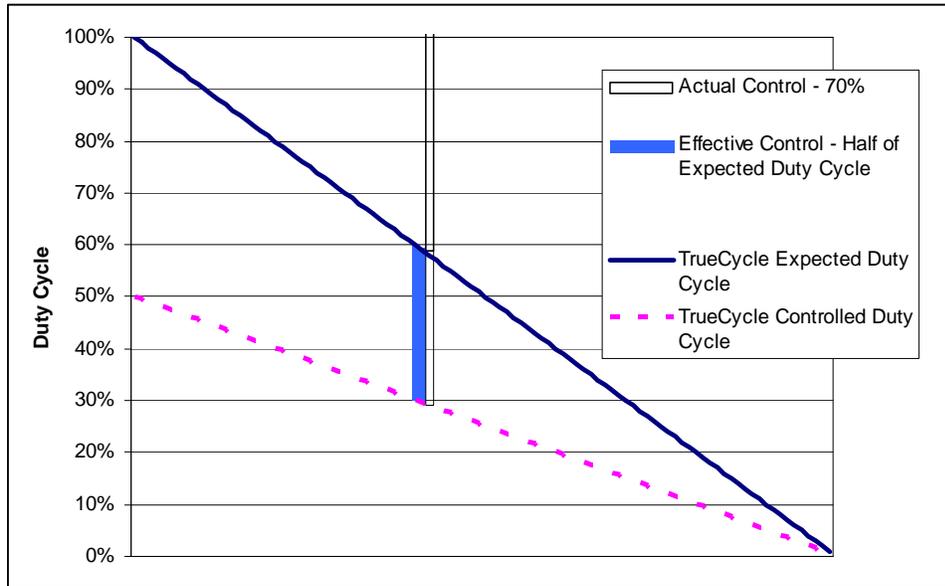
The program-level effectiveness of traditional switches is limited by AC units that are not to running full time. Oversized units and mild weather both cause “natural”, un-controlled run-times to be less than full time, thus lowering the avoided run-time. Oversizing, in particular, is an issue that undermines load reduction even under the extreme conditions that motivate a DLC program event. This limitation to the effectiveness of traditional switches has led to a new generation of adaptive switches designed to adjust to the natural duty cycle of the AC unit. The PG&E SmartAC Program uses Cannon’s TrueCycle technology adaptive switch.

The adaptive switches address the traditional switch limitation by “learning” the run-time behavior of the unit. The program administrator or system operator chooses learning days that have the characteristics of potential event days. The observed amount of run-time on these learning days provides an estimate of expected run-time on an event day. With this information the switch determines a more appropriate control level.

For the purpose of explanation, it is easiest to think in terms of duty-cycle rather than run-time. Duty-cycle is in the percentage of an hour during which the compressor is on. Control is the percentage of an hour during which duty cycle is restricted from running. The traditional switch maintains a constant, “actual” 50 percent level of control. That actual 50 percent control is only an effective 50 percent of the natural duty cycle if the unit is running at 100 percent. TrueCycle is designed to approximate an effective control level of 50 percent (in the case of the SmartAC Program) of the natural duty cycle even if the natural duty cycle drops below 100 percent.

Figure 3-1 provides a visual representation for an expected natural duty cycle of 60 percent. If the expected duty cycle is 60 percent, then the TrueCycle process will cut that in half to a 30 percent duty cycle. It will limit the unit to a maximum 30 percent duty cycle by enforcing an actual control level of 70 percent. As the expected duty cycle decreases from 100 percent (left to right) the actual control increases from 50 percent to 100 percent to maintain an effective control of 50 percent of the expected duty cycle.

**Figure 3-1
Actual vs. Effective Control for Expected Duty Cycle of 60 Percent**



The success of the adaptive switch in overcoming the limitations of the traditional switch relies on the estimate of expected duty cycle. At the beginning of the cooling season, or any time the process fails, the default expect duty cycle for any hour is 100 percent. Under these default conditions, the TrueCycle switch control is identical with the traditional switch. As learning days are identified, an average duty cycle is calculated that includes the observed duty cycle from the learning day. Cannon indicates that they generally use a weight of one eighth for a single learning day. If fewer than eight learning days have been identified then the remaining days included in the mean calculation are assumed to be at the default of 100 percent. Using this approach a rolling estimate of expected duty cycle for each hour is maintained for each AC unit. These estimates are supposed to represent expected duty cycle under the extreme conditions likely for program event day. Cannon also indicates that the estimated expected duty cycle may be adjusted to pre-event duty cycle levels though the mechanics of this adjustment are not provided.

The TrueCycle adaptive technology has a number of implications for the estimation of load impacts. Most importantly, traditional switch performance is the lower bound for the adaptive switches. The adaptive load reduction can only improve on that offered by the traditional switches. In this respect, the TrueCycle technology has the potential to address the limitations of traditional switches with essentially no risk of decreasing load reduction.

The actual effectiveness of the TrueCycle technology in estimating expected duty cycle is much more difficult to determine. The number and choice of learning days drives the estimate. Extreme conditions are experienced only infrequently so the data on AC unit usage under extreme conditions is, in fact, sparse. To the extent less extreme days are included in the calculation, the adaptive algorithm has the potential to under-estimate an AC unit's extreme day duty cycle. In this case, the participant could potentially face an effective control level greater than 50 percent. Under this scenario, the overall load reduction is increased, but participants may be experience greater control-related discomfort.

Finally, it is important to reiterate two switch characteristics in the context of the discussion of the effectiveness of load reduction. Switches are essentially invisible to the participant. A light on the switch indicates control mode, but few participants will notice this as the switch is installed at the AC unit. Furthermore, as the switch is designed to provide regular periods of control, it also provides regular periods of cooling. In addition, when the compressor is turned off, the circulation fan still operates⁵. The operating characteristics of the switch technology facilitate the “invisibility” of the process even under event conditions.

On the other hand, switches do not directly control indoor temperature. The increase of indoor temperatures will be a function the controlled unit run-time, outdoor temperature and household characteristics. Across a population, the range of indoor temperature increase will vary, with some participants at the high end experiencing increases of greater than five degrees.

3.1.3 Thermostats

Thermostat control technology directly controls household temperature⁶. When activated to event mode, thermostats increase the temperature to which the house is cooled. The unit may turn off if already in cooling mode. If the unit is already off, it may remain off for a longer period so as to allow the thermostat to reach the new, higher setpoint temperature. Using the thermostat setpoint as the focus of control puts the premium on controlling the increase in participant indoor temperatures. No participant should experience an indoor temperature increase greater than the setpoint. In theory, increasing the thermostat setpoint equitably distributes temperature increase across the participating population regardless of house and AC unit characteristics.

As indicated, the direct control of thermostat setpoint has an indirect effect on AC energy usage. How an AC unit responds to the setpoint increase will be a function of the pre-event cooling regime, the cycling schedule of the AC unit, house-specific characteristics affecting the rate of indoor heat gain, and the amount of setpoint increase. The most common scenario involves the AC unit turning off (or staying off) until the indoor temperature reaches the level of the higher setpoint. For this period, while the house warms to the new setpoint equilibrium, program-related savings are 100 percent of the pre-program usage. Once the new equilibrium is reached, the AC unit returns to cycling behavior necessary to maintain cooling at this higher setpoint. As AC usage is fundamentally a function of differential with outside temperature, usage at the new setpoint will be reduced relative to pre-event usage levels. However, compared to the interim period of readjustment, the energy usage at this new equilibrium temperature will be greater.

⁵ Depending on the characteristics of the house and the placement of the unit, this may or may not have a positive effect on indoor temperature. For example, if air circulates through uninsulated pipes in very hot attics, it may increase indoor temperature more than if the fan was off at the same time as the compressor.

⁶ The thermostats employed by the SmartAC Program have traditional switch capabilities. In the future, the thermostats will have TrueCycle adaptive algorithm capability. For the purpose of this evaluation, however, we focus on the setpoint-related capabilities of the thermostat.

When thermostats were first deployed for DLC programs, the most common control strategy was a single setpoint increase of three, four or five degrees Fahrenheit. A setpoint increase of this sort has two important implications for DLC programs. First, the load reduction is greater during the first part of the event. During the initial period of readjustment the majority of units turn off as the indoor temperature slowly increases to the new setpoint. This high level of load reduction is maintained until units reach their new equilibrium temperature. From a system load reduction perspective, this inability to maintain a constant level of load reduction can be seen as a limitation.

Second, a single block setpoint increase has implications of a parallel nature for participants' perceptions of comfort. The period of readjustment to the new setpoint is experienced by the inhabitant as a period of no cooling. For what could be a substantial period of time, there is no blowing of the circulation system, no experience of cooling air. From a participant comfort perspective, the block set point increase also has limitations.

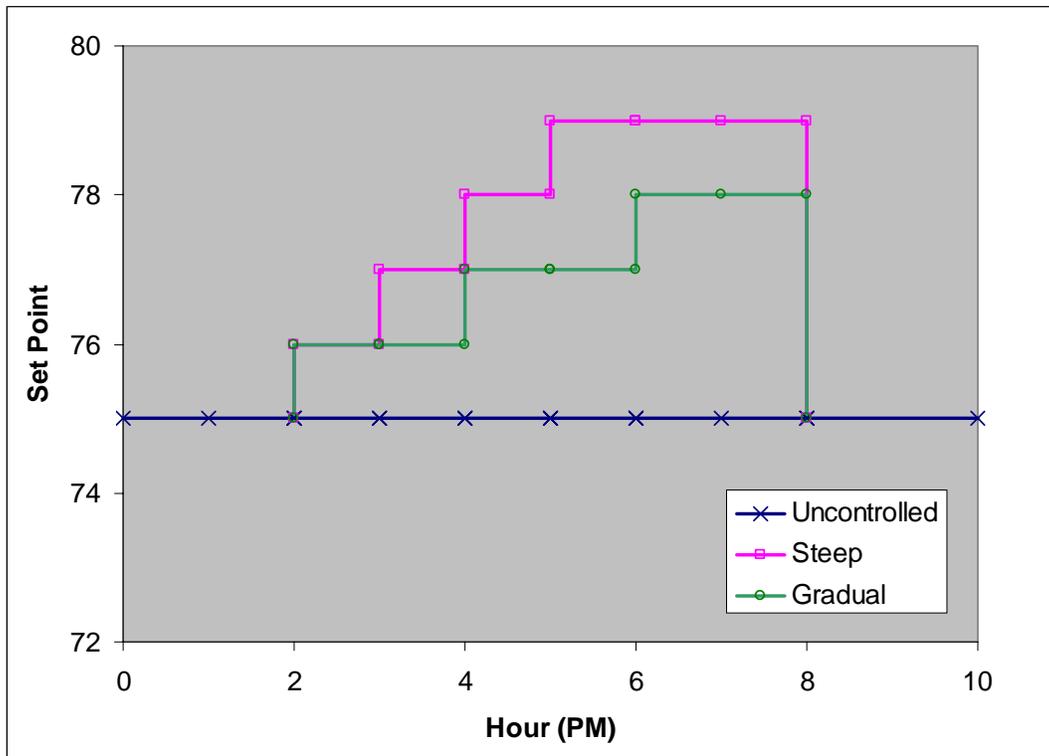
Utilities have experimented with ramped control strategies to combat these twin limitations of the block setpoint increase. Multiple increases of a degree or two should, in theory, mitigate both problematic aspects of the single block setpoint increase. Instead of experiencing the load reduction and resulting lack of cooling in one long period at the beginning of the event, the periods of temperature equilibrium readjustment are shorter and spread out through the event.

The SmartAC program chose to use a ramping strategy for all program thermostat participants if there were full-program events during summer 2007. Once the program was deployed the ultimate strategy chosen imposed an increase of 1°F at the beginning of the first, third, and fifth hours of the event⁷. We refer to this strategy as the “gradual” strategy. For the purpose of testing the ramping strategy concept, PG&E identified a second, more aggressive ramping strategy to be tested only on the meter sample which we refer to as the “steep” strategy.

Figure 3-2 diagrams the two strategies applied in an alternating pattern to the two halves of the thermostat sample during the summer of 2007. The “gradual” strategy is represented by three equal steps of one degree over two hours. The “steep” strategy increases the setpoint 1°F at the beginning of each of the first four hours. After the fourth increase, the “steep” strategy stays at a 4°F increase for the duration of the event. All else being equal, the “steep” strategy ought to provide greater impacts, at least during certain hours, than the “gradual” strategy.

⁷ Prior to program deployment, PG&E had anticipated that the “steep” strategy would consist of setting back thermostats 4°F on the first hour, and that the “gradual” strategy would consist of increasing 1°F at the beginning of each of the first four hours. Prior to initiating the EM&V events, PG&E changed the strategies to what is described in this section.

Figure 3-2
Comparison of Ramping Strategies for 2 pm Start Time



The goal of both ramping strategies is to combat the limitations of a single block setpoint increase: front-weighted impacts and long periods with no cooling activity. This evaluation directly assesses the effect of the ramping strategy on load reduction.

4. Data and Methodology

This section summarizes the data and methods used to conduct this evaluation.

4.1 Data

This section describes the data used to complete this load impact evaluation, and their sources. The data includes:

- Program enrollment data
- Load data collected specifically for this EM&V process
- Billing data
- Weather data from PG&E
- Process evaluation data collected specifically for this EM&V process.

4.1.1 Program Enrollment Data

Program tracking data files were provided by PG&E's implementation contractor. Site, work order, device and dropout data were received in separate datasets. Datasets were received in a rolling fashion for the purpose of selecting the EM&V sample. The impact evaluation made use of tracking data provided in November, for the purposes of developing weights that were specific to each DLC event.

Important fields from the tracking data include:

- Site data
 - Customer ID
 - Structure size
 - Structure age
- Work order data
 - Date installed
- Device data
 - Type of device
 - Tonnage
- Dropout Data
 - Date removed

Unit tonnage is central to the load impact analysis presented in this report. Unit tonnage is relatively easy to collect from the unit nameplate and is strongly correlated with unit connected load. Impact results are calculated as impact per ton. As of August of 2007, 86 percent of the AC unit records in the tracking data included unit tonnage.

Structure size and structure age were important for the imputation process used to fill the missing unit tonnage data. Both fields were present for about 98 percent of the units in the tracking data. Average unit tonnage was computed for each combination of structure size and structure age categories. This average was used to impute tonnage for records with missing tonnage values.

Date installed and removed determined the rolling program population total for any given event day. Units were assumed to be active a day after installation.

SEER would have been extremely useful information to have. Unfortunately, SEER is frequently not available on the unit nameplate. Alternative approaches to finding SEER like model number look-ups are time-consuming and have their own limitations. It is usual for DLC program evaluation to lack SEER data from program participants.

Table 4-1 provides a summary of the program tracking data as of August 31st, 2007. The table provides the number of sites and units, and the average tons across the full population, control technology sub-groups and single and multiple units.

**Table 4-1
SmartAC Program Population Statistics as of August 31st, 2007**

Category	Program			Thermostat Control			Switch Control		
	All	Single Unit	Multi Unit	All	Single Unit	Multi Unit	All	Single Unit	Multi Unit
Participant Sites	8,193	7,605	588	2,445	2,282	163	5,748	5,323	425
Units	8,843	7,605	1,238	2,610	2,282	328	6,233	5,323	910
Average Tons	3.3	3.3	3.3	3.2	3.2	3.2	3.4	3.4	3.3

4.1.2 Load Data

Sample design took place in early summer during a period of rapid population increase. The final sample was created using program tracking data as of June 11th, 2007. The sample, by necessity, reflected the program as it existed at that time. The sample was pulled with the intent of metering all units at the selected sites.

This section discusses the original sample design utilized to deploy the loggers used to collect load data, and the re-stratification performed for this analysis. The re-stratification supports the ratio-estimation approach used to aggregate unit-level results to program and control-type results. This approach addresses a concern of this load impact study -potential bias due to changes in the composition of the population that occurred after the sample had been selected.

4.1.2.1 Sample Design and Re-Stratification

The original sample design had eight strata. This design was based on type of device, tons per site and whether or not multiple ACs are present at the site. The tracking data variable that captured unit tonnage had approximately 17 percent missing data. For the sample design, these tonnage missing values were

imputed through the analysis of billing data supplied by PG&E. The sample design that was used to implement the logger sample is provided in Table 4-2.

**Table 4-2
Original Sample Design For Logger Data Collection**

Stratum	Type Of Device	Total Tons From All Units	Multiple AC units on site (1=Yes)	Program Participants as of 06/11/2007	Design Sample Size	Design Number of Loggers	Number of Metered Homes	Number of Metered AC Units
1	PCT	<4	0	483	93	93	88	90
2	PCT	<4	1	6	3	6	2	4
3	PCT	>=4	0	148	37	37	39	39
4	PCT	>=4	1	34	17	34	17	37
5	Switch	<4	0	1,404	77	77	72	73
6	Switch	<4	1	21	5	10	5	10
7	Switch	>=4	0	637	54	54	53	53
8	Switch	>=4	1	123	17	34	21	46
Totals				2,856	303	345	297	352

For the impact evaluation it was desirable to estimate unit impact as a function of tons. With this in mind, KEMA conducted an additional effort to fill the missing tonnage data in a more informative manner. This process is discussed below. Because the impact evaluation results use the tonnage variable, it was necessary to re-stratify the sample using the new imputations. As a result of the re-stratification, strata 2 and 6, which had multiple units but small tonnage, became impractically small. To address this problem, strata 2 and 6 were combined with strata 4 and 8, respectively. This collapsed stratification, within each device type, preserved the stratification of single unit sites into less than 4 tons (“small”) or 4 tons or more (“large”), but groups all multiple unit sites into the same stratum.

The ultimate effect of the re-stratification is small. Table 4-3 illustrates the collapsed strata with the missing tons imputation method that was employed for the sample design.

**Table 4-3
Original Sample Design Collapsed to Six Strata**

Stratum	Type Of Device	Total Tons From All Units	Multiple AC units on site (1=Yes)	Program Participants as of 06/11/2007	Percent of Total Population by Device Type	Design Sample Size	Design Number of Loggers	Number of Metered Homes	Number of Metered AC Units
1	PCT	<4	0	483	72%	93	93	88	90
3	PCT	>=4	0	148	22%	37	37	39	39
4	PCT		1	40	6%	20	40	19	41
5	Switch	<4	0	1404	64%	77	77	72	73
7	Switch	>=4	0	637	29%	54	54	53	53
8	Switch		1	144	7%	22	44	26	56
Totals*				2856		303	345	297	352

*Original sample participant count was determined by when the file was created, June 11th, 2007.

Table 4-4 shows the population and sample after re-stratification on the improved tonnage variable. In terms of the program population, the re-stratified sample allocation is almost identical to the original sample. For the re-stratification process we used tracking data from the end of the cooling season. It includes both the population used for the original sample and the final population needed to characterize the program through the summer. We used installed date to recreate a population similar to the population used for the original sample. Table 4-4 illustrates the effect of re-stratification on the original population counts, by limiting the counts to units installed by June 4th, 2007

**Table 4-4
Sample Design with Re-Stratification Based on Improved Tonnage Variable**

Stratum	Type Of Device	Total Tons From All Units	Multiple AC units on site (1=Yes)	Program Participants as of 06/11/2007	Percent of Total Population by Device Type	Design Number of Loggers	Design Number of Loggers	Final Sample Size	Final Number of Loggers
1	PCT	<4	0	501	73%	110	110	88	88
3	PCT	>=4	0	139	20%	30	30	39	39
4	PCT		1	44	6%	10	19	21	45
5	Switch	<4	0	1431	64%	96	96	74	74
7	Switch	>=4	0	633	28%	43	43	48	48
8	Switch		1	169	8%	11	23	27	59
Totals*				2917		300	321	297	353

*Re-stratified sample participant count uses an installed date of June 4th, 2007. Final number of loggers reflects an additional logger installed at a multi-unit site.

Table 4-5 presents the final population numbers for the 2007 cooling season. The last event of the season took place on September 26th, 2007. At that time, there were about 9,600 participants enrolled in the program. Interestingly, though the population increased more than threefold between June and September, and thermostat adopters increased from about one fourth to about one third of the program's participants, the percent of the device population represented by each strata is identical when rounded.

**Table 4-5
SmartAC Population as of September 26th, 2007, Re-stratified**

Stratum	Type Of Device	Total Tons From All Units	Multiple AC units on site (1=Yes)	Program Participants as of 09/26/2007	Percent of Total Population by Device Type
1	PCT	<4	0	2275	73%
3	PCT	>=4	0	625	20%
4	PCT		1	208	7%
5	Switch	<4	0	4170	64%
7	Switch	>=4	0	1850	28%
8	Switch		1	503	8%
Totals*				9631	

4.1.2.2 Onsite Data Collection - Metering

To enable the impact estimate analysis methodologies that KEMA planned to use for this evaluation, KEMA installed a combination of one-minute and 15-minute loggers on all units at the sites selected for the meter sample. All sites using switch control were monitored using one-minute loggers. Sites with thermostat controls were logged with a mix of 15- and one-minute loggers.

The loggers used were the HOBO Energy Logger ProTM (for one-minute data) and the DENT DATA^{pro}TM Data Logger (for 15-minute data.) Both loggers used a current transformer (CT), installed around a single leg of an air-conditioning unit, to monitor the voltage of the electromagnetic field produced by an alternating current, and were programmed to convert that voltage reading into amps. Four different sized CTs were used for this study: 20, 70, 100, or 150 amp. The data that was stored in both loggers for each reading during the metering period captured the following information:

- Date
- End Time
- Average Amp
- Plot Number

The DATApro™ loggers were programmed to capture 15-minute interval data. The loggers took instantaneous amp readings every minute and recorded the average of those readings at the end of fifteenth minute.

The HOBO Energy Logger Pro™ is a modular, reconfigurable data logging system which was combined with the S-FS-TRMSA FlexSmart TRMS Module to record an instantaneous amp reading every minute.

4.1.2.3 Data Cleaning

Data cleaning was performed at two levels.

- **Logger level quality control**-- Determine that the logger operated correctly and recorded air conditioner compressor load levels.
- **Interval level** -- Identify and remove intervals that are deemed unreliable.

4.1.2.3.1 Logger Quality Control

The primary data cleaning decision was determining whether the logger represented valid readings of compressor load for the site and the AC unit. To answer this question, a number of different pieces of information were considered:

- **Device issues.** Some loggers failed in the field and others failed in the data download process. In both cases, data was unavailable.
- **AC energy use signature.** The majority of AC compressors have only two modes, on or off. Logger data from AC compressors (both one- and fifteen-minute) have a distinct signature. Connected load, the load when the compressor is on, is a function of ambient temperature. When logger data is plotted with respect to hourly temperature, a clear trend is evident in compressor logger data. Each logger was plotted and determined to have a signature that reflected air conditioner usage.
- **Comparison of aggregated logger load data to monthly billing data.** Loggers with sparse data may represent a failed logger or a valid AC unit that was not used very frequently. To address this concern we compared logger usage data with billing data provided by PG&E. For each site, we compared site AC usage over the billing period to premise kWh from billing data. AC usage generally drives summer energy usage and a clear parallel trend exists between whole premise and the AC-only data. For the majority of sites a clear visual relationship exists between billing data and aggregated logger data. In instances where billing data indicated weather-sensitive load and the aggregated logger data did not show air conditioner usage, the loggers were considered faulty.

To include a logger in the second level of cleaning, the logger data needed to satisfy expectations on all of the three counts described above. Table 4-6 provides the results from the initial logger quality control check. The table classifies reasons for removal from this analysis into three categories.

**Table 4-6
Logger Disposition**

Strata	Reason for Removal from Analysis			Total
	Billing Data Comparison	Failed logger	No Tracking	
1	8	5	0	13
3	7	2	0	9
4	4	3	0	7
5	8	4	0	12
7	3	1	0	4
8	3	6	1	10
No Strata	0	0	1	1
Total	33	21	2	56

Failed loggers include loggers that were:

- Broken or damaged when collected by the technician,
- Found disconnected, removed or otherwise not installed properly, or
- Failed to download the recorded data properly.

The loggers with no tracking data were loggers installed on units that did not correspond to those identified in the program tracking data.

The largest number of failed loggers is classified under the Billing Data Comparison field. The billing data comparison became the final hurdle questionable loggers needed to clear. Loggers with sparse data or with data lacking the compressor load signature might still accurately record units that were hardly used through the summer. On the other hand, the data on the logger could be bad data. If the billing data provided clear evidence of a summer cooling load then the logger was removed and classified under this category. Many loggers with little to no usage passed this test because the billing data did not provide evidence of cooling load.

4.1.2.3.2 Interval Cleaning

The cleaning of individual loggers started with the trimming of the partial day at the beginning of the logging period and all data recorded after November 1st when KEMA started collecting the loggers from the field. This trimming removed almost all anomalous intervals. In the fifteen-minute data, only three loggers remained with a handful of negative or extremely high readings. The intervals were set to missing and the logger-level QC process reconsidered to confirm that despite the anomalous loads removed, these loggers contained otherwise valid data.

In addition to trimming the data, the data was checked for missing intervals and intervals not consistent with the air conditioning usage signature discussed above. There were no missing data intervals present other than those created by the cleaning process. The data do include a very small number of intervals that show higher load than is consistent with the energy use signature. These intervals are few in number and are consistent with an instantaneous read that coincided with the initial energy draw at the beginning of a duty cycle. While these few artificially high intervals are easily detected in the interval data plotted by hourly temperature, artificially low readings are difficult to impossible to discern, especially in the fifteen minute logger data. In the interest of not biasing cleaned data, these few possibly questionable intervals were not removed.

4.1.2.3.3 Final Logger Disposition

Table 4-7 provides the final strata counts of metered homes and AC units.

**Table 4-7
Final, Post-Cleaning Number of Homes and AC Units by Strata**

Stratum	Type Of Device	Total Tons From All Units	Multiple AC units on site	Number of Metered Homes	Number of Metered AC Units
1	PCT	<4	0	76	76
3	PCT	>=4	0	30	30
4	PCT		1	17	35
5	Switch	<4	0	62	62
7	Switch	>=4	0	44	44
8	Switch		1	25	50
Totals				254	297

Table 4-8 provides the counts of loggers used to estimate the results reported in this analysis. There was sample attrition due to two program participants leaving the program. A third sample participant opted out of the sample but remained in the program.

**Table 4-8
Sample Unit Counts by Event for All Treatment Sub-groups**

Event Date	Event Temp (DAT)	Program Logger Count	Switch Count	Thermostat Count		
				All	Gradual Ramp	Steep Ramp
7/12/2007	71	297	156	141	141	
7/17/2007	73	297	156	141	69	72
7/23/2007	81	297	156	141	72	69
7/26/2007	75	297	156	141	69	72
7/27/2007	76	297	156	141	72	69
8/1/2007	77	297	156	141	69	72
8/9/2007	75	297	156	140	71	69
8/10/2007	78	297	156	140	69	71
8/21/2007	84	297	156	140	71	69
8/22/2007	83	297	156	140	69	71
8/28/2007	84	295	155	139	70	69
8/30/2007	87	295	155	139	69	70
8/31/2007	88	295	155	139	70	69
9/10/2007	74	295	155	139	69	70
9/26/2007	74	295	155	139	139	

4.1.2.4 Conversion of Logger Readings to kW

Both the one-minute and 15-minute loggers recorded amps. The one-minute loggers recorded an instantaneous amp reading every minute. The 15-minute loggers measured instantaneous amp readings every minute, and recorded the average of those readings every 15 minutes. For the 15-minute data produced from the one-minute loggers, we calculated the average amps over the 15 minutes, consistent with the process internal to the 15-minute loggers. One-minute intervals were aggregated to 15-minute data prior to interval cleaning. Thus, 15-minute data from the one-minute loggers was produced and cleaned identically with the 15-minute logger data. The available one-minute interval data for use at the one-minute level was cleaned at that level.

The amperage was converted to kW using the voltage levels measured at the units. Voltage measures were attempted at all units. Where they were not successful, we substituted the average measured voltage calculated across all measured units. Among the units where voltage was successfully measured, there was relatively little variation across measured voltage.

4.1.3 Weather Data

For the purpose of the impact analysis, PG&E provided observations of dry bulb temperature and relative humidity in half-hour intervals for three weather stations covering the initial phase of the SmartAC Program (Stockton and surrounding areas) as of summer of 2007, for the period from January 1 through December 31, 2007.

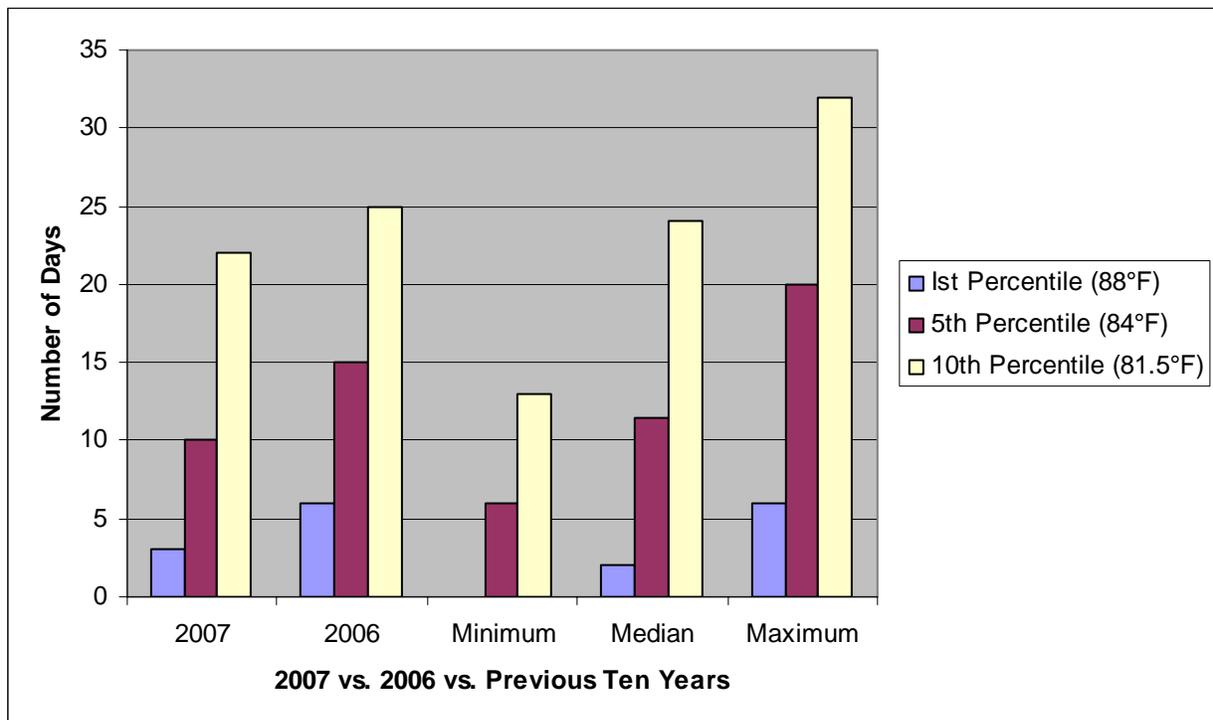
The load and duty cycle modeling used daily average temperature (DAT) as the weather variable in all models. The daily average temperature is calculated as:

$$\text{DAT} = (\text{Maximum Temperature} + \text{Minimum Temperature})/2$$

In addition, PG&E provided historical weather data for Stockton that included observations of daily average temperature for the period from May 1 through October 31 for the years 1983-2006. These data established percentiles cut-offs to identify the one, five and ten percent hottest days across the 25 years. The first percentile contains days with daily average temperature above 87.5° F. The fifth percentile contains days with daily average temperature above 83.1° F. These two cut-offs in particular are of interest for the SmartAC program because they represent the days when the program is likely to operate for system-related reasons, and when meaningful impacts are generated.

The historical weather data allow us to characterize the summer of 2007 in relation to past years, particularly the hot summer of 2006. Based on counts of days above the three thresholds, 2007 was a typical summer compared to the ten previous summers. Figure 4-1 shows the range and median number of days at each threshold. The 2007 counts are slightly above the median for the 1st percentile days but below for the 5th and 10th percentiles.

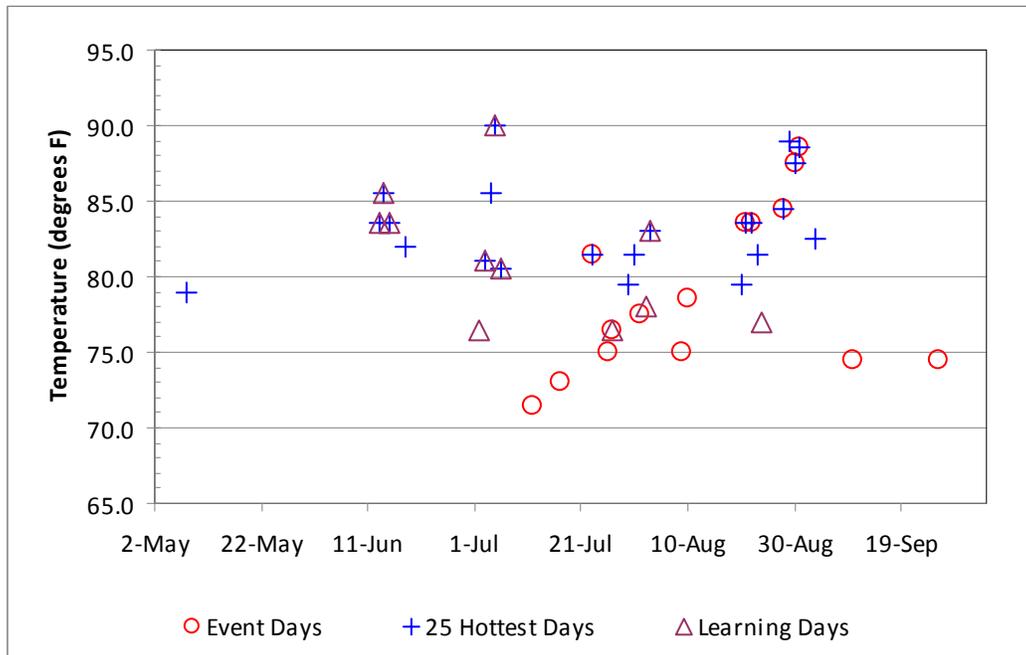
Figure 4-1
Comparison of Summer 2007 to Previous Decade, Count of Days in 1st, 5th and 10th Percentiles



The summer of 2006, along with 1998, had the highest number of day in the first percentile temperatures at or above 88° F. The summer of 2007 only had half as many of these most extreme days compared to 2006. The summer of 2007 only had two thirds of the 5th percentile days compared to 2006.

Figure 4-2 shows the hottest days from the summer of 2007 with the 15 event days and the 11 learning days⁸. This figure shows the large number of events called during relatively mild conditions. It also shows that there were non-event days across the spectrum of temperature. Non-event days in all temperature ranges are essential for modeling purposes. Finally, Figure 4-2 shows that a number of quite mild days were included as learning days for the TrueCycle technology.

Figure 4-2
Summer 2007 Hottest Days, Event Days and Learning Days



4.2 Overview of Analysis Methodologies

For this evaluation, two different modeling approaches were used to estimate expected AC load on event days. We used a basic load model to estimate load for units with the thermostat control technology. In addition, a duty cycle model was used to estimate load for the switch participants.

4.2.1 Load Model

4.2.1.1 Basic Model

The basic kW model estimates load as a function of dry bulb temperature in the form of average daily heating or cooling degree-days. Using hour-specific dummy variables, the intercept and both degree-day measures are included in the model on an hour-specific basis. This means that each of the 24 hourly load

⁸ “Learning days” are days when the Program operators send a signal to the SmartAC Switches, to record the air conditioner’s activity on that day. The devices use this information to implement adaptive load curtailments on event days.

indicator variables for each day are regressed against an hour-specific intercept term and degree-day term. The resulting parameters, though based on only a single daily temperature measure, provide an hourly estimate of load as a function of weather.

Equation 1

$$L_{ihd} = \alpha_{ih} + \beta_{Hih} H_d(\tau_{Hi}) + \beta_{Cih} C_d(\tau_{Ci}) + \varepsilon_{ihd}$$

where

- L_{ihd} = sum of 15-minute interval AC consumption at hour h of day d for unit i ;
- $H_d(\tau_{Hi})$ = heating degree-days at the heating base temperature τ_{Hi} for unit i on day d , based on daily average temperature;
- $C_d(\tau_{Ci})$ = cooling degree-days at the cooling base temperature τ_{Ci} for unit i , on day d , based on daily average temperature;
- ε_{ihd} = regression residual;
- $\alpha_{ih}, \beta_{Hih}, \beta_{Cih}$ = coefficients determined by the regression; and
- τ_{Hi}, τ_{Ci} = base temperatures determined by choice of the optimal regression.

The degree-day variables are calculated as:

$$C_d(\tau_{Ci}) = \max((F_d - \tau_{Ci}), 0)$$

$$H_d(\tau_{Hi}) = \max((\tau_{Hi} - F_d), 0)$$

The model is fit separately for each AC unit across a range of heating and cooling degree day bases. The ideal cooling base temperature is the minimum ambient temperature at which AC use begins and below which there tends to be no AC load. The ideal heating base temperature, where relevant, is the maximum ambient temperature above which there is no heating load. Base temperatures vary across premises because individuals' indoor temperature preferences vary and because homes vary in their physical properties that influence indoor temperature.

We also estimate the model with no heating parameters and no heating or cooling parameters as well. We then compare all the different model specifications using an appropriate F-test to choose the combination with the best explanatory power for each premise.

4.2.1.2 kW Model Load Estimates

The optimal kW model for each unit includes a set of estimated parameters. Depending on the optimal model chosen, the model may or may not include heating and cooling parameters. The most common optimal model including only base and cooling parameters is provided in Equation 2.

Equation 2

$$\hat{L}_{ihd} = \hat{\alpha}_{ih} + \hat{\beta}_{Cih} C_d(\hat{\tau}_{Ci})$$

Where the hat variables on the right hand side represent estimated parameters from the regressions and \hat{L}_{ihd} is the estimated load for unit i in hour h on day d .

The basic weather normalization model estimates a base load as well as heating and cooling parameters. Where AC load is the only dependent variable being modeled, we would expect this base load to be zero unless the AC unit is a heat pump, or there is some ongoing, low-level load used by the condenser. In instances where the weather normalization model produced base load parameters that were, in aggregate, negative, we set the base load parameters to zero. This restricted weather normalization model is identical to the basic weather normalization model shown in Equation 2 except that it lacks the α_{jh} parameters.

4.2.2 Duty Cycle Modeling

Duty cycle modeling directly estimates a unit’s duty cycle as a function of some weather variable. To do this kind of modeling, duty cycle must be derived from the logger data. At the same time the unit’s connected load (full load kW) must also be derived so the duty cycle can be converted back to kW. The resulting duty cycle-based modeling is amenable to estimating the impacts of program controls that function on the unit’s duty cycle.

4.2.2.1 Deriving Duty Cycle and Connected Load

Deriving duty cycle and connected load amounts to careful cleaning of one-minute logger data based on knowledge of how air conditioner units work.

One-speed air conditioner compressors are either on or off. If the logger is truly recording only compressor load then each one-minute interval of kW is either zero or the full load kW for that unit. Duty cycle, the percent of an hour the compressor is running, is easily calculated by counting the number of minutes in the hour during which the logger was at full load kW.

Each unit’s non-zero full load kW can be expressed as a linear function of hourly temperature:

Equation 3

$$CL_{idh} = \alpha + \beta f_{dh}$$

where

- CL_{idh} = Connected load for unit i for day d , hour h
- f_{dh} = Dry bulb temperature for day d and hour h
- α, β = Estimated parameters

Equation 4 provides the equation for estimated connected load. The fit values of this equation represent the connected load for that unit across the full range of cooling temperatures. In general, connected load increases approximately one percent per degree Fahrenheit increase.

Equation 4

$$\hat{C}L_{idh} = \hat{\alpha} + \hat{\beta}f_{id}$$

$\hat{C}L_{idh}$ = Connect load for unit i for day d , hour h

$\hat{\alpha}, \hat{\beta}$ = Estimated parameters

Plotting and performing this simple regression of logger data reveals that air conditioner data does, for the majority of intervals, conform to this simple structure. The process also reveals readings that do not conform to this structure. Two kinds of data commonly fall in this category:

- Fans and other auxiliary devices picked up by the logger,
- High kW readings likely measuring the onrush of current as the compressor starts up.

Data points representing each of these data issues are easily identified. Non-compressor loads picked up by the logger are generally small and have a distinct structure with regards to temperature. The high data points are rare and randomly situated outside of the realistic bounds of compressor load. Because there is so much data supporting the connected load relationship with temperature, both of these undesired readings are easily removed from the connected load regressions.

Once the modeled definition of connected load is derived, a decision is made as to how these readings are characterized with regard to connected load. Fans and other auxiliary loads below the level of connected load indicate the compressor is off. The rare, random high loads are explained as a current onrush and deemed to represent a running compressor.

In combination the duty cycle and connect load derived from the one-minute kW intervals fully represent the load used by a unit during any hour.

4.2.2.2 Estimating Expected Duty Cycle

The duty cycle component of the analysis consisted of estimating the expected duty cycle under a range of times and conditions. Duty cycle is a function of cooling demand. Cooling demand is, in turn, a function of the difference between the actual and desired internal temperatures of the cooled space. Actual temperature reflects the heat building up in the house over the course of the day. The desired internal household temperature is completely determined by human behavior as indicated by thermostat setpoints. Hourly duty cycle will reflect the combination of these two factors.

To start, a simple linear relationship between duty cycle and temperature is hypothesized. Daily average temperature provides the best measure of the daily heat load of a house. In this case, human behavior is too varied to account for directly in a statistical model. However, by estimating separate models for each hour between 12 noon and 10 PM, we can account for general patterns of behavior that might have changed over the hours of the afternoon.

The duty cycle analysis recognizes that the relationship between desired cooling and temperature is essentially linear, but the cooling provided cannot be less than zero or more than 100 percent of the unit’s capacity. We therefore use a statistical method that assumes an underlying unconstrained linear relationship of desired cooling, which would be negative (desired heating) at low temperatures and exceed 100 percent capacity at high temperatures. The theoretical unconstrained duty cycle also has random variation around the line. The observed duty cycle data are assumed to be the result of this underlying linear relationship with random noise, constrained to a minimum of zero and maximum of 100 percent runtime.

The model form that has this structure is a Tobit model. Details of the model structure and its application to the duty cycle analysis are provided in Appendix A.

Figure 4-3
Duty Cycle Model Schematic

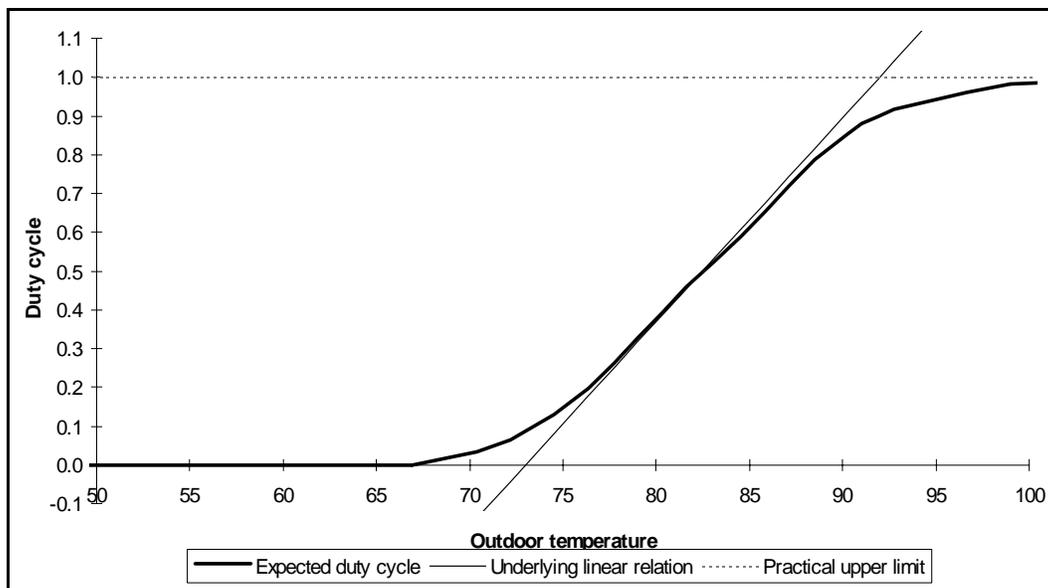


Figure 4-3 illustrates the Tobit duty cycle model. The straight diagonal line represents the assumed underlying linear relationship between, in this case, expected unconstrained duty cycle and daily average temperature. The S-shaped curve shows how the model estimates the resulting constrained relationship between expected duty cycle and temperature. For any given temperature, the values on the S-shaped curve indicate the expected duty cycle.

4.2.2.3 Duty Cycle Model Load Estimates

Load estimates for the duty cycle approach are calculated by combining duty cycle and connected load, as shown in Equation 5:

Equation 5

$$\hat{L}_{ihd} = \hat{D}C_{ihd} \hat{C}L_{ihd}$$

4.2.3 Choice of Weather Data

Daily average temperature is the primary weather variable used in the kW model. Daily average temperature captures the range of temperatures experienced by the house through the day. Houses are temperature integrators. That is, they heat up and cool down more slowly than ambient temperatures. The previous night's temperature is an important indicator of how much of the previous day's heat the house shed during the night. This information is equally as important as maximum temperature in the estimation of cooling load.

KEMA also included humidity and lagged temperature variables in the modeling efforts using a temperature-humidity index and a weighted three day average temperature. Our conclusions in the modeling process were that humidity plays a less significant role in driving cooling load in the Stockton area. Lagged temperature variables, while more promising, did not improve estimates for enough units to justify the added complexity lagged temperatures bring to projected impacts.

4.2.4 Impact Estimates

The unit level savings regardless of estimation approach are calculated per Equation 6:

Equation 6

$$S_{ihd} = \hat{L}_{ihd} - L_{ihd}$$

Both the kW model and the duty cycle model estimate load on an hourly basis as indicated in the equation above. The load data, however, were available on a quarter-hour basis. Kilowatt-hour impacts for each quarter-hour interval were calculated analogously to the hourly equations indicated above. For the quarter-hourly estimates, the load in each time increment was estimated used the load model coefficients for the hour that included that increment.

Impacts were also calculated at the unit level for each hour and for the fully active period of the event. The first half hour of each event was left out of event level calculation because units are randomly activated through that first half hour and thus all units are not yet active. The event level results do include through the official end of each event as that is the point at which units start returning to normal operation.

Estimates of snapback include the 90 minutes starting 30 minutes after the official end of the event. Once again, the random nature of each unit's start and end means that the first 30 minutes after the official end of the event will include a mix of controlled and un-controlled units.

4.2.5 Projected Impact Estimates for General Conditions

4.2.5.1 Load Model Projections

Projected impact estimates are derived from the unit-level kW models. For any hour, daily average temperature and setpoint increase, the estimated impact is:

$$\text{Equation 7}$$

$$\bar{S}_{if\Delta h} = \hat{\beta}_{Cih} C_{f\Delta}(\hat{\tau}_{Ci})$$

where

$\bar{S}_{if\Delta h}$ = Estimated impact for unit i for hour h at temperature f for temperature differential Δ , and

$C_{f\Delta}(\hat{\tau}_{Ci})$ = The difference in cooling degree-days given the cooling base temperature τ_{Hi} for unit i for hour h at temperature f for temperature differential Δ .

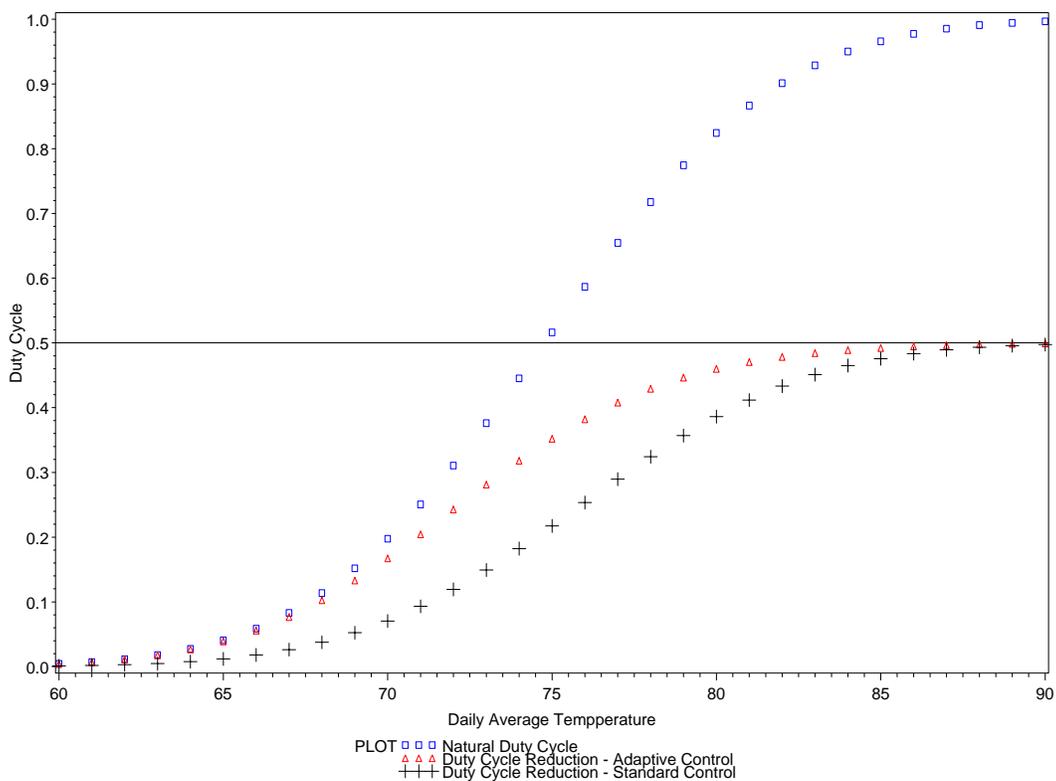
If the daily average temperature is above the cooling degree day base then there will be impacts. Once the daily average temperature is more than the temperature differential above the degree day base then the projected impact becomes constant.

4.2.5.2 Duty Cycle Model Projections

The duty cycle reduction due to cycling control is the difference between the uncontrolled duty cycle and actual control level, if the uncontrolled duty cycle would be above that actual control level. Even when the expected duty cycle is below actual control level, there is some probability that the uncontrolled duty cycle would exceed actual control level because of random variation around the S-shaped expected duty cycle curve. Thus, the expected duty cycle reduction is positive even when the expected duty cycle itself is below actual control level.

The fitted Tobit model for each unit and time period allows estimation not only of the expected duty cycle but also of the expected duty cycle reduction as a function of temperature and actual control level. Figure 4-4 shows a plot with the natural duty cycle and duty cycle reduction representing two different levels of control. Standard control is the flat, 50 percent control that has been commonly used by switch programs. The adaptive control effectively controls duty cycle to half the level of the natural duty cycle. The duty cycles reduction associated with these two levels of control provide an approximate bound to the duty cycle reduction expected with the adaptive algorithm used by the adaptive switches installed for SmartAC. Details of the specific calculations used are found in the Appendix.

Figure 4-4
Natural Duty Cycle with Duty Cycle reduction for 50 Percent and Adaptive Control



4.2.6 Ratio Estimation

Unit-level results must be expanded to represent the full program population. The ratio estimation approach we used estimates impact per ton. Ratio estimation has a number of advantages over a non-ratio based approach: Improved precision, lower bias due to changes in population composition and the requested direct estimate of impact per ton.

The ratio estimator is calculated as shown in Equation 8.

Equation 8

$$\frac{\bar{S}_{hd}}{\bar{T}} = \frac{\sum_{k=1}^6 \sum_{c=1}^{n_k} \sum_{i=1}^{m_c} w_{kci} S_{kcihd}}{\sum_{k=1}^6 \sum_{c=1}^{n_k} \sum_{i=1}^{m_c} w_{kci} T_{kci}}$$

where

- S_{kcihd} = gross demand impact of the for CAC i in hour h on day d ,
- T_{kci} = tons for CAC i in cluster c and strata k ,
- w_{kci} = sampling weight for impact of CAC i , identified as CAC k in cluster c of stratum q ,
- k = stratum number with a total of n_q clusters,
- c = cluster number within stratum k ,
- i = CAC unit number within cluster c of stratum k ,
- n_k = number of clusters in stratum k ,
- m_c = number of units in cluster c ,
- \bar{S}_{hd} = gross demand impact of the average CAC in the program population in hour h at daily average temperature d ,

With sampling weights calculated as:

Equation 9

$$w_{kci} = (N_k / n_k)(M_{kc} / m_{kc})$$

- n_k = total number of clusters (sites) in meter sample, strata k ,
- N_k = total number of clusters (sites) in program population, strata k ,
- m_{kc} = Number of CACs in meter sample, strata k and cluster c ,
- M_{kc} = Number of CACs in program population, strata k and cluster c ,

The ratio estimator result provides the estimate of unit load reduction per ton. The final estimate of unit load reduction for the events of the summer are calculated as shown in Equation 10:

Equation 10

$$\overline{S}_{hd} = \frac{\overline{S}_{hd}}{\overline{T}} \overline{T}^*$$

Where \overline{T}^* is the average tons for the population on the day of the event. Using the ratio estimator addresses concerns about bias resulting from sample drift as the program grows. If the character of the population changes with respect to the size of the unit, then that change will be accounted for by using the average tons from the day of the event.

In these equations h=hour but it also could be 15-minute intervals or the full event period. Unit level data is aggregated to the desired level before combining across units with the ratio estimation process.

4.3 Drivers of the Impact Results

The results reported in Section 5 reflect the estimated per-unit impacts of the SmartAC program from a number of different perspectives. While temperature is the fundamental driver of program impacts, a wide array of factors interact with temperature to determine how much impact the program will have under different conditions. Three participation-related factors directly affect the potential impact of demand response programs: (1) non-use of the air conditioner, (2) device or signal failure and (3) event opt-out. In addition, the different control strategies for thermostats also played a major role in the impact outcomes.

4.3.1 Participation-Related Factors

Non-use of air conditioners – To register reductions in usage an air conditioner has to be used. Air conditioner usage is ultimately determined by customer behavior. Even at the hottest temperatures, there will be customers who use their air conditioner minimally or not at all. Units that are never turned on and, thus, generate no savings, are included and reduce the impact estimates.

Device or signal failure – This category encompasses any failure in the process of controlling an individual unit. This could include the paging signal not being received by the switch or thermostat. It could also be a failure of the control device to successfully control the unit despite receiving the broadcast signal. In some older demand response programs, extremely high rates of control device/signal failure have been identified. In a new program such as SmartAC, it is reasonable to expect relatively low failure rates.

Control device and signal failure are accounted for in the estimated impacts reported above. It is assumed that this kind of failure is present in the meter sample at level representative of failure in the population as a whole. Units experiencing this kind of failure will not be controlled and thus will decrease the overall estimated impact.

Under certain circumstances, it is possible to improve the accuracy of impact estimates by treating units with some form of failure as a separate group. This approach does not affect the estimated mean impact but it

can decrease the variation around that estimate. This results in lower standard errors and tighter confidence intervals. Some demand control programs use two-way technology that provides a confirmation of receipt of the paging signal. This confirmation is available for all participants in the program population, not just the metered sample. This population level signal failure data can be used to account for that failure without contributing to the error of the estimated impact. The SmartAC program does not use technology with two-way capability. The expected low failure rate in a program such as SmartAC would limit the increased accuracy gained with this approach.

Event opt-out – Many demand response programs offer participants the option of over-riding (opting out of) program control when they so choose. The over-ride option can be an important feature for selling demand response programs to customers. There are a wide range of over-ride options used by different programs. The initiation of the over-ride can take place at the thermostat or can involve internet or telephone access. In some programs, over-ride options include a monetary disincentive. When over-ride has a low cost in terms of money or hassle, hot day over-ride rates can be as high as 50 percent. When the over-ride option plays this kind of role in a program, it is essential to model its effect on the program directly.

Event opt-outs for the SmartAC program were very few. Over all fifteen EM&V events, there were only 13 instances of over-rides by 7 meter sample participants. Two of these participants eventually dropped out of the SmartAC program. The opt-out option played such a small role in the 2007 EM&V SmartAC events that it was not possible to assess its effect as a separate driver of program savings. The impact results presented in this report are based on all participants including those participants who opted-out and, thus, take into consideration opt out. The projected results are based solely on the underlying unit-specific models and do not account for potential future opt outs.

As program participation extends beyond “early adopters”, and PG&E changes the solicitation and post-installation materials provided to SmartAC participants to increase awareness of opt-out options, it is possible there will be an increase in the number of participants who choose to opt out of future program events.

5. Impact Evaluation Results

This section presents the impact evaluation results. We report results separately for each of the fifteen event days during the 2007 cooling season. The results include

- Overall program impact estimates.
- Impact per unit for each control technology,
- Hourly results, overall and by control technology
- Comparison between thermostat ramping strategies
- Impact by site and unit characteristics
- Projected impacts
- Connected load, and
- Snapback.

As discussed in the methodology section, different modeling approaches were applied to the two control technologies. For the switch-based controls, we modeled duty-cycle (run-time) and connected load (kW when running) at the unit level, directly from one-minute interval data. The duty cycle model facilitates modeling switch-based control strategies like the adaptive algorithm used by the SmartAC program. For thermostat controls, we modeled air conditioner kW as a function of weather variables at the unit level. The models provide a unit-specific estimate of load across the full range of temperature conditions.

We use the terms “estimated impact” or “observed impact” to refer to the impact calculated as the difference between a baseline (calculated with our regression models) and the observed load at the time of an event. We use the term “projected impact” to refer to the impact predicted by our regression models under specific conditions.

5.1 Event Level Results

Table 5-1 provides load impact estimates for the full SmartAC program, in order of observed descending daily average temperature. The estimated demand reduction on the hottest event day generated an average of 1.21 kilowatts of load savings for each participating unit. The five hottest event days generated impacts with statistical precision of 11 to 17 percent at the 90 percent confidence level. For August 31st, the statistical precision of 11 percent at the 90 percent confidence level indicates that if repeated samples were pulled from the SmartAC population, 90 percent of the time the estimated per unit impact would fall between +/- 11 percent of the this estimate of 1.21 kW per unit. A common goal in impact evaluations is 90/10 precision or 10 percent precision at the 90 percent confidence level. This level of precision is particularly challenging for run-time oriented end-use like air conditioners. A precision of greater than 100 percent means the per unit estimate cannot be statistically distinguished from zero.

**Table 5-1
Program Impact Results, Average Per-Unit kW per Event(*)**

Event Date	Event Day Daily Average Temperature	Event Start	Event Duration	Per Unit Impact (kW)	Participating Units	Program Wide Impact (MW)	Confidence Interval Lower Bound (MW)	Precision at 90 Percent Confidence (90/xx)
8/31/2007	88	2 PM	5	1.21	8,843	10.7	9.5	11%
8/30/2007	87	3 PM	4	1.02	8,809	9.0	7.8	13%
8/28/2007	84	3 PM	4	0.85	8,690	7.4	6.6	12%
8/21/2007	84	2 PM	5	0.71	8,306	5.9	5.1	13%
8/22/2007	83	2 PM	5	0.56	8,391	4.7	3.9	17%
7/23/2007	81	12 PM	4	0.18	6,566	1.2	0.6	44%
8/10/2007	78	2 PM	4	0.35	7,613	2.6	2.1	20%
8/1/2007	77	2 PM	4	0.04	7,049	0.3	-0.2	> 100%
7/27/2007	76	2 PM	4	0.06	6,825	0.4	0.0	97%
8/9/2007	75	2 PM	4	0.12	7,523	0.9	0.5	43%
7/26/2007	75	2 PM	4	-0.01	6,757	-0.1	-0.5	> 100%
9/26/2007	74	2 PM	5	0.30	10,344	3.1	2.6	16%
9/10/2007	74	4 PM	3	0.06	9,297	0.5	0.0	> 100%
7/17/2007	73	12 PM	5	-0.02	6,215	-0.2	-0.4	> 100%
7/12/2007	71	2:30 PM	3.5	-0.01	5,949	0.0	-0.3	> 100%

(*) Event Averages do not include the first half hour of the event period when all participants are not yet activated.

Impact results are estimated using the difference between site-level estimated load or duty cycle and actual load or duty cycle. On mild days when there is little cooling there may be effectively no savings. In these cases, the per-unit impact represents the model error relative to observed load. When this is the case small negative impact results are possible.

Table 5-2 presents results by control device. On the hottest event day of the summer, August 31, 2007, the average impact for the switch control is 50 percent higher than the average impact for the thermostats. Over the five hottest days, the average impact for the switch controls is almost twice the average impact for the thermostats. These hot day impact differences between thermostats and switches are all statistically significant.

**Table 5-2
Thermostat and Switch Impact results, Average Per-Unit kW per Event**

Event Date	Event Day Daily Average Temperature	Thermostat Participants			Switch Participants			Per-Unit Impact Difference Statistically Significant
		Per-Unit Impact (kW)	Standard Error	Precision at 90 Percent Confidence (90/xx)	Per-Unit Impact (kW)	Standard Error	Precision at 90 Percent Confidence (90/xx)	
8/31/2007	88	0.89	0.11	21%	1.34	0.10	12%	X
8/30/2007	87	0.48	0.11	38%	1.24	0.11	14%	X
8/28/2007	84	0.45	0.09	33%	1.02	0.08	12%	X
8/21/2007	84	0.51	0.08	27%	0.79	0.07	15%	X
8/22/2007	83	0.27	0.09	52%	0.68	0.07	18%	X
7/23/2007	81	0.10	0.09	> 100%	0.20	0.06	46%	
8/10/2007	78	0.25	0.06	42%	0.38	0.05	22%	X
8/1/2007	77	-0.02	0.08	> 100%	0.07	0.05	> 100%	
7/27/2007	76	0.00	0.07	> 100%	0.09	0.04	86%	
8/9/2007	75	0.20	0.05	38%	0.08	0.04	75%	X
7/26/2007	75	0.03	0.07	> 100%	-0.02	0.04	> 100%	
9/26/2007	74	0.34	0.05	25%	0.28	0.04	21%	
9/10/2007	74	0.08	0.06	> 100%	0.05	0.05	> 100%	
7/17/2007	73	0.00	0.04	> 100%	-0.03	0.03	> 100%	
7/12/2007	71	-0.09	0.04	79%	0.02	0.03	> 100%	X

(*) Thermostat control results are a combination of the two ramping strategies used by the program.

The sample was designed to achieve 10 percent precision at the 90 percent confidence level for overall program results. Results at the control technology level were not expected to attain this level of precision. The results in Table 5-2 illustrate that variation across units with thermostat control is generally greater than the variation across units with switch control.

5.2 Impact at Time of System Peak

PG&E experienced its 2007 system peak on August 29th at 5pm. This day was not a SmartAC EM&V event day⁹. August 28th, 30th and 31st were event days. Because of the proximity of these event days to the day of the system peak, these days provide the best indication of potential program impact at the time of the system peak.

⁹ SmartAC EM&V event days were scheduled on the day prior to the event. This is addressed in the Process Report that is part of the SmartAC evaluation project.

Table 5-3 and Figure 5-1 provide hourly impact estimates for the SmartAC program for the three days around the day of the system peak. These are also the three hottest event days of the summer of 2007. The system peak took place at 5 pm. The lowest average kW reported at this time was 0.78 kW on August 28th. The highest average kW was 1.23 kW on the 31st. These hours have precision levels of 12 to 15 percent at ninety percent confidence.

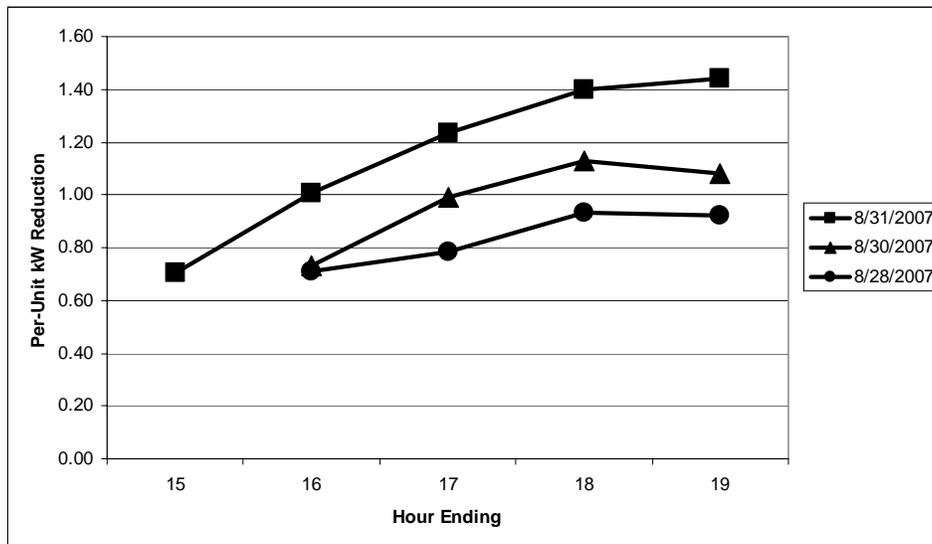
**Table 5-3
Program Impact results, Average Per-Unit kW per Hour on Hottest Event Days**

Event Date	Hour Ending	Per Unit Impact (kW)	Participating Units	Program Wide Impact (MW)	90 Percent Confidence Interval Lower Bound (MW)	Precision at 90 Percent Confidence (90/xx)
8/31/2007 Temperature: Average = 88 Maximum = 98	3 PM	0.70	8,843	6.2	4.9	21%
	4 PM	1.01	8,843	8.9	7.6	14%
	5 PM	1.23	8,843	10.9	9.6	12%
	6 PM	1.40	8,843	12.3	11.0	11%
	7 PM	1.44	8,843	12.7	11.6	9%
8/30/2007 Temperature: Average = 87 Maximum = 100	4 PM	0.73	8,809	6.4	5.2	20%
	5 PM	0.99	8,809	8.7	7.4	15%
	6 PM	1.13	8,809	9.9	8.6	13%
	7 PM	1.08	8,809	9.5	8.3	13%
8/28/2007 Temperature: Average = 84 Maximum = 99	4 PM	0.71	8,690	6.2	5.1	17%
	5 PM	0.78	8,690	6.8	5.8	14%
	6 PM	0.93	8,690	8.1	7.1	12%
	7 PM	0.92	8,690	8.0	7.1	12%

(*) Participating units are number of residential units enrolled in the program on the day of the event

Figure 5-1 shows the increase in impacts across the days and the trend of impacts across the hours on each of the three event days. August 29th, the day of the system peak, had a daily average temperature of 89°F for the EM&V sample sites. The maximum temperature for that day was 103°F, between the hours of 4 and 6 pm. Air conditioner usage and measured impacts generally increase as heat waves extend to multiple days. This is clearly the case for this four day period. This trend would point to a system peak impact estimate falling between the estimates of the 28th and 30th. The 5 PM midpoint between these two days is 0.88.

Figure 5-1
Program Impact Results, Average Per-Unit kW Reduction per Hour on the
Three Hottest Days of the Load Impact Evaluation



On the other hand, August 29th, the day of the peak and the second of the four-day heat wave, was the hottest. It was a half a degree hotter in terms of daily average temperature than August 31st, the next hottest of the four and last in the four day sequence. If temperature alone was considered the driver then the average kW estimates of 1.23 and 1.40 kW from August 31st would be the best proxy values for the system peak estimated impact.

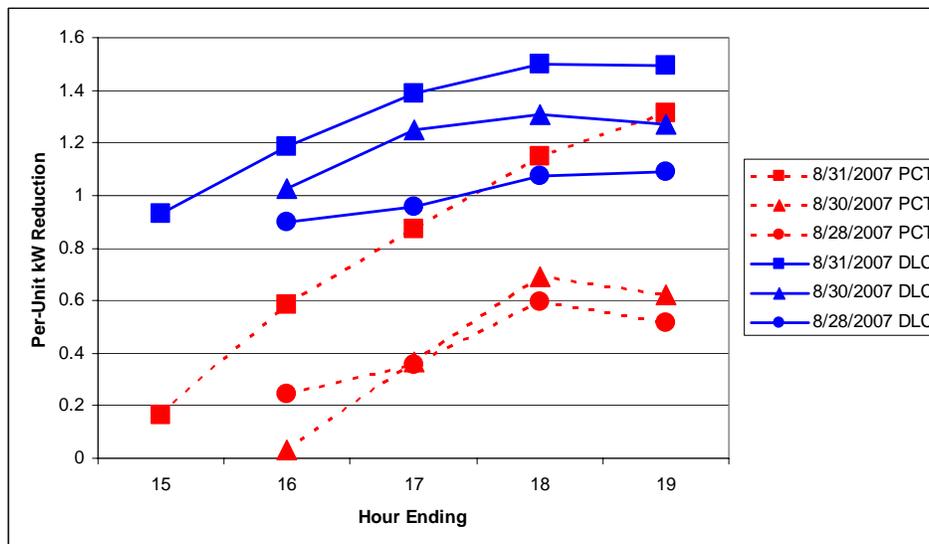
Table 5-4 and Figure 5-2 provide results by control device. These data reinforce the event-level results indicating higher impacts from the switches.

Table 5-4
Thermostat and Switch Impact Results, Average Per-Unit kW per Hour on Hottest Event Days

Event Date	Hour Ending	Thermostat Participants			Switch Participants			Per-Unit Impact Difference Statistically Significant
		Per-Unit Impact (kW)	Standard Error	Precision at 90 Percent Confidence (90/xx)	Per-Unit Impact (kW)	Standard Error	Precision at 90 Percent Confidence (90/xx)	
8/31/2007 Temperature: Average = 88 Maximum = 98	3 PM	0.17	0.13	> 100%	0.93	0.11	20%	X
	4 PM	0.58	0.12	34%	1.19	0.11	16%	X
	5 PM	0.87	0.14	27%	1.39	0.11	13%	X
	6 PM	1.15	0.13	19%	1.50	0.11	13%	X
	7 PM	1.32	0.13	16%	1.49	0.10	11%	X
8/30/2007 Temperature: Average = 87 Maximum = 100	4 PM	0.03	0.13	> 100%	1.02	0.11	18%	X
	5 PM	0.37	0.13	60%	1.25	0.11	15%	X
	6 PM	0.69	0.14	33%	1.31	0.12	15%	X
	7 PM	0.62	0.13	34%	1.27	0.10	14%	X
8/28/2007 Temperature: Average = 84 Maximum = 99	4 PM	0.24	0.12	83%	0.90	0.09	17%	X
	5 PM	0.36	0.12	54%	0.96	0.08	14%	X
	6 PM	0.59	0.11	31%	1.07	0.09	14%	X
	7 PM	0.52	0.11	36%	1.09	0.08	12%	X

Of greater interest, Figure 5-2 shows that the two devices have different impact trends through the hours of the events. The thermostat impacts are much lower than the switch impacts during the early hours. By the last hour of the events, the difference is reduced, and in the case of August 31st, substantially so.

Figure 5-2
Thermostat and Switch Impact Results, Average Per-Unit kW per Hour on Hottest Days



This result is unexpected given the ramping strategies employed by the program. Generally thermostat impacts for un-ramped controls are “front-loaded”. The largest impacts come early in the event when all units are off as the house temperatures increase to the new set point. Once units turn back on to maintain the house temperature at the higher set point, the impacts moderate. Ramping strategies are employed to address this issue. The ramping smoothes out the initial high impacts across the whole event period. The evidence here indicates that an initial one degree increase might be too small to generate early impacts. It also suggests that other factors, such as home occupancy at the time of the event, may be different for each of the technology groups.

5.3 Comparison of Thermostat Ramping Strategies

Thermostat sample participants were randomly assigned to two groups, and stayed in the same group throughout the summer. The two thermostat control strategies were applied to these two thermostat groups alternately. This alternation was designed to control for differences in the thermostat subgroups in the overall load modeling. Even with the random assignment, it is evident from these results that one of the subgroups consistently generated greater impacts than the other. For instance, all of the days when the two strategies are statistically different are days when the group with greater impacts received the “steep” ramp strategy.

Table 5-5 provides the results for the two ramp strategies. As expected the “steep” strategy has higher impacts for all of the five hottest days. On only two of these days, however, is the difference between the two strategies statistically significant.

**Table 5-5
Gradual and Steep Thermostat Ramp Strategy results, Average Per-Unit kW per Event**

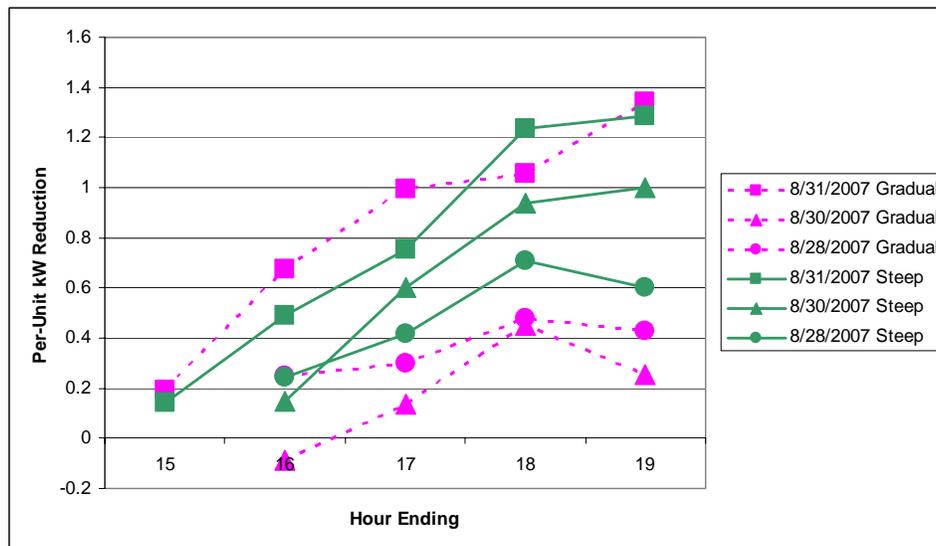
Event Date	Event Day Daily Average Temperature	Gradual Ramp			Steep Ramp			Per-Unit Impact Difference Statistically Significant
		Per- Unit Impact (kW)	Standard Error	Precision at 90 Percent Confidence (90/XX)	Per- Unit Impact (kW)	Standard Error	Precision at 90 Percent Confidence (90/XX)	
8/31/2007	88	0.93	0.15	27%	0.85	0.17	33%	
8/30/2007	87	0.23	0.15	> 100%	0.75	0.17	38%	X
8/28/2007	84	0.38	0.11	47%	0.53	0.15	47%	
8/21/2007	84	0.39	0.11	47%	0.64	0.12	33%	
8/22/2007	83	0.10	0.12	> 100%	0.44	0.12	46%	X
7/23/2007	81	0.24	0.11	79%	-0.04	0.13	> 100%	
8/10/2007	78	0.23	0.08	58%	0.27	0.10	61%	
8/1/2007	77	-0.24	0.13	88%	0.19	0.10	86%	X
7/27/2007	76	0.03	0.09	> 100%	-0.03	0.10	> 100%	
8/9/2007	75	0.13	0.07	92%	0.28	0.06	38%	
7/26/2007	75	-0.01	0.10	> 100%	0.07	0.10	> 100%	
9/26/2007	74	0.34	0.05	25%				
9/10/2007	74	-0.08	0.09	> 100%	0.25	0.08	56%	X
7/17/2007	73	0.01	0.05	> 100%	-0.01	0.07	> 100%	
7/12/2007	71	-0.09	0.04	79%				

Table 5-6 and Figure 5-3 focus on the hourly results for the three hottest event days of the 2007 cooling season. These results repeat the findings presented above at the event level. Only during the August 30th event did the steep ramp provide statistically significantly higher impact results, and then only in the last three hours. As stated above, this appears to reflect differences in the two thermostat subgroups in addition to the difference in ramping strategies.

Table 5-6
Gradual and Steep Thermostat Ramp Strategy results, Average Per-Unit kW per Hour on Hottest Event Days

Event Date	Hour Ending	Gradual Ramp			Steep Ramp			Per-Unit Impact Difference Statistically Significant
		Per- Unit Impact (kW)	Standard Error	Precision at 90 Percent Confidence (90/xx)	Per- Unit Impact (kW)	Standard Error	Precision at 90 Percent Confidence (90/xx)	
8/31/2007 Temperature: Average = 88 Maximum = 98	3 PM	0.19	0.18	> 100%	0.14	0.19	> 100%	
	4 PM	0.68	0.16	40%	0.49	0.18	62%	
	5 PM	0.99	0.19	31%	0.75	0.22	48%	
	6 PM	1.06	0.17	27%	1.24	0.21	28%	
	7 PM	1.34	0.17	21%	1.29	0.19	25%	
8/30/2007 Temperature: Average = 87 Maximum = 100	4 PM	-0.09	0.19	> 100%	0.15	0.17	> 100%	
	5 PM	0.13	0.18	> 100%	0.60	0.19	53%	X
	6 PM	0.45	0.17	64%	0.94	0.22	38%	X
	7 PM	0.25	0.19	> 100%	1.00	0.17	29%	X
8/28/2007 Temperature: Average = 84 Maximum = 99	4 PM	0.25	0.17	> 100%	0.24	0.18	> 100%	
	5 PM	0.30	0.15	82%	0.42	0.19	74%	
	6 PM	0.48	0.14	49%	0.71	0.17	41%	
	7 PM	0.43	0.12	45%	0.60	0.20	54%	

Figure 5-3
Gradual and Steep Thermostat Ramp Strategy results, Average Per-Unit kW per Hour on Hottest Event Days



5.4 Load Reduction by Site and Unit Characteristics

This section examines in more detail the difference in load reduction across the two control device populations. These differences could reflect the differences between these two self-selected populations (since Program participants had a choice between each of the two devices) or differences in the control technologies.

The following figures compare the thermostat and switch load reduction results across two important unit-level categories:

- Unit size of less than four tons vs. unit size of four tons or more, and
- Single vs. multiple AC units

5.4.1 Large vs. Small AC Units

This evaluation looked at impact per ton relative to unit size for the two different control technology groups. The results are consistent with the expected interaction between control technology and oversized AC units.

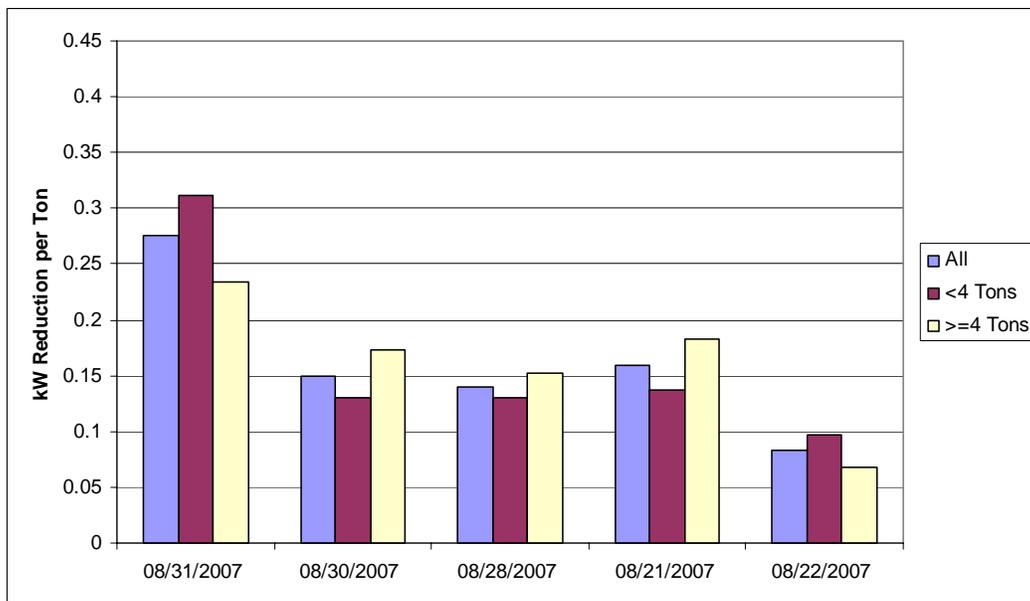
If all units are properly sized we expect similar usage per ton regardless of unit size. All else being equal, impact per ton should also be similar across different sized units. The different control technologies limit usage in different ways, but with right sizing that limitation should be consistent across unit size.

This consistency across size changes if oversizing is present to a higher degree in larger units. An adaptive switch will, at best, reduce usage by a consistent 50 percent regardless of sizing. However, an oversized unit by definition has lower usage per ton, thus lower impact per ton. If more oversized units are present in the large AC category then we would expect per ton impacts to be lower for that category.

To the contrary, within the range of normal usage the thermostat control has a constant effect regardless of unit size.¹⁰ Figure 5-4 and Figure 5-5 compare two size groups for each kind of control device on the five hottest days of summer 2007. If oversizing is present in the population, it is more likely to be in the larger size group.

¹⁰ At extremes this will not be the case. If a unit is running constantly and not keeping the house cool then an increase in setpoint may not lower usage at all. This has the effect of lowering impacts for undersized units.

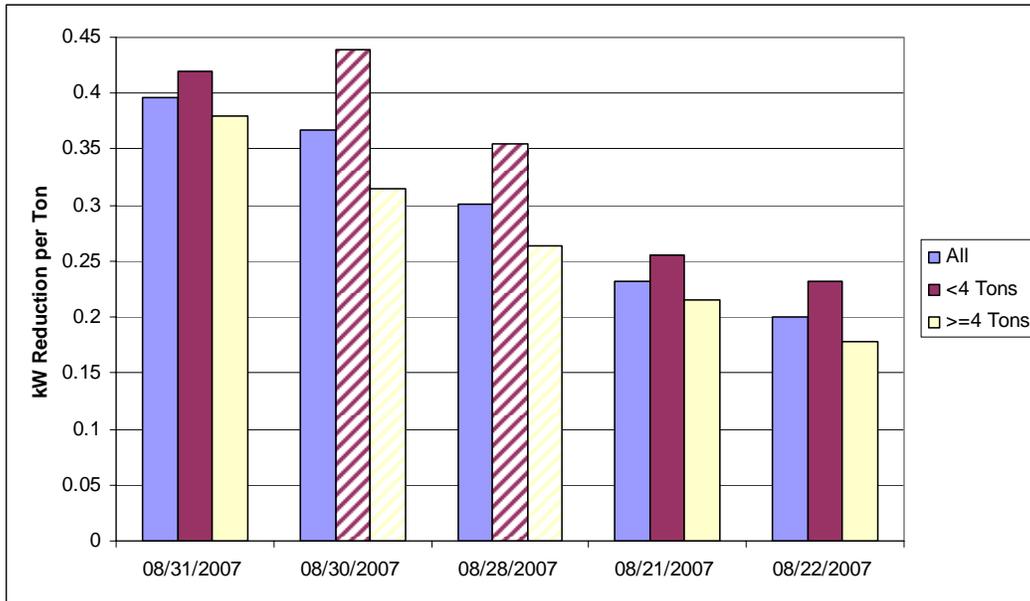
Figure 5-4
Thermostat Control Load Reduction per Ton, Large vs. Small Units



In Figure 5-4, for units with thermostat controls, there is no consistent pattern of per-ton load reduction comparing large and small units. None of the differences are statistically significant. This is consistent with the constant thermostat control effect.

In Figure 5-5, for units with switch controls, small unit per ton load reduction is higher than large unit load reduction for all five days. The difference is statistically significant for two of the five days. Assuming a presence of oversize units in the population, this is consistent with the expected sensitivity of switches to oversizing.

Figure 5-5
Switch Control Load Reduction per Ton, Large vs. Small Units



* Dates with stripes, subgroup results are statistically different

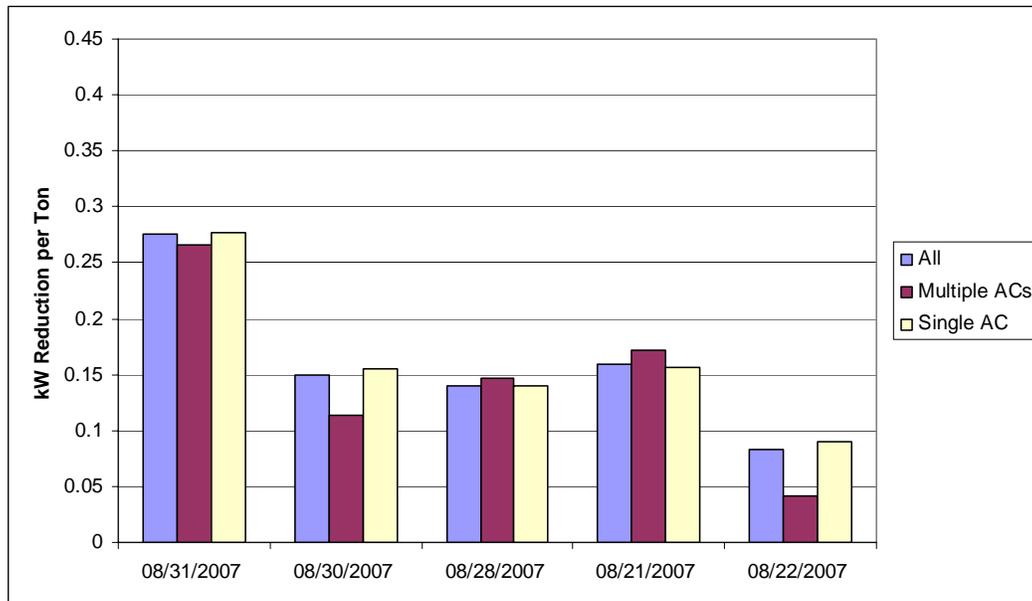
These findings are only suggestive. The differences between the two control technology groups may extend to rates of oversizing. Assuming oversizing is equally distributed, and also that the larger size group is more likely to contain oversized units, then these results support our hypothesis. More importantly, these results point to the need for further exploration of the effect of oversizing on control method efficacy.

5.4.2 Single vs. Multiple AC Units

The presence of multiple air conditioning units is also a potential driver of load reduction. The general concern is that households with multiple AC will deliver lower load reduction per ton. This can be explained by secondary units that are programmed to run during limited hours or because control periods are not synchronized. We find evidence of this for the SmartAC program, but it is only consistent and statistically significant for the households with switch controls.

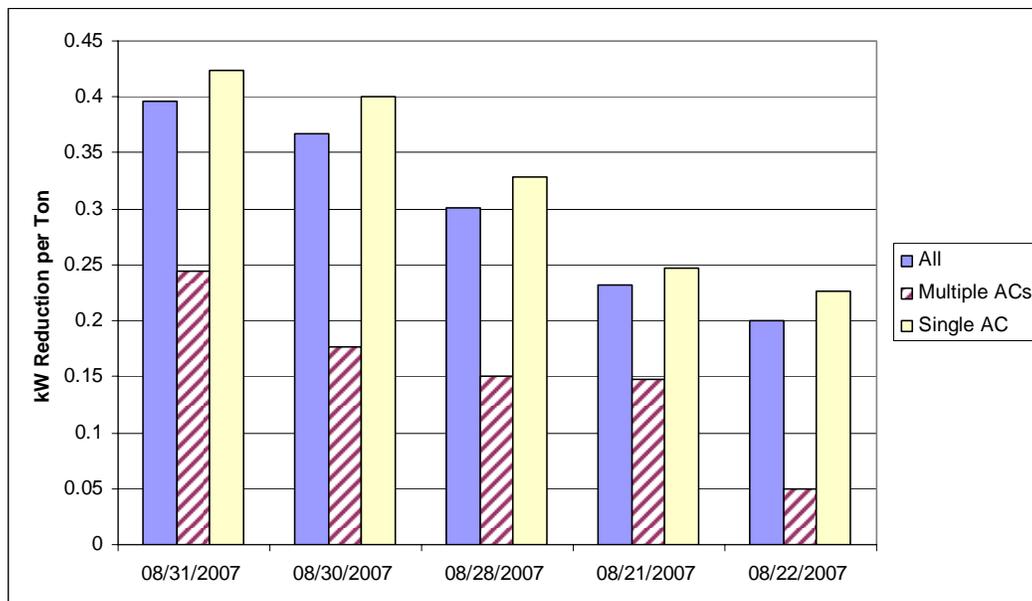
Figure 5-6 and Figure 5-7 show per-ton load reduction for the five hottest event days of the summer for single and multiple unit households. For thermostat households, multiple unit household load reduction is similar to single unit household load reduction. On all five of these days the difference between single and multiple AC households is not statistically significant.

Figure 5-6
Thermostat Control Load Reduction per Ton, Single vs. Multiple Units



For households with switch control, those with multiple units supply approximately half the load reduction per ton of the single unit household units. On all five days, the difference is statistically different.

Figure 5-7
Switch Control Load Reduction per Ton, Single vs. Multiple Units



* Dates with stripes, subgroup results are statistically different

These differences across the two self-selected populations are dramatic. Both cooling behavior and control technology are at work here. Control technology effects alone are unlikely to cause this substantial difference.

5.5 Projections for 2008

One of the purposes of this evaluation was to establish predicted load reductions across a range of temperatures, durations and start-times. These projections may be used for planning purposes for summer 2008. These per-unit projections multiplied by the number of units active for summer 2008 events give estimates of program load reductions.

The control strategies used for both switches and thermostats create challenges for calculating projected impacts. Both thermostat ramping strategies as well as the adaptable switching algorithm introduce control elements that are difficult to generalize. The ramping strategies, for instance, are specific to both the hour of the day as well as the hour of the event. The projected impact during the hour ending at 5 pm will differ depending on when the event started as well as the daily average temperature and ramping strategy. The TrueCycle technology is even more difficult to account for as it is a unit-specific learning algorithm that adapts to unit usage as the summer progresses.

5.5.1 Switch Projections for 2008

The TrueCycle adaptive algorithm makes impact projections for units with switch controls more challenging. Projections always take models based on specific data (in this case, summer 2007) and project results across a range of scenarios. With the TrueCycle adaptive algorithm, the impacts will vary

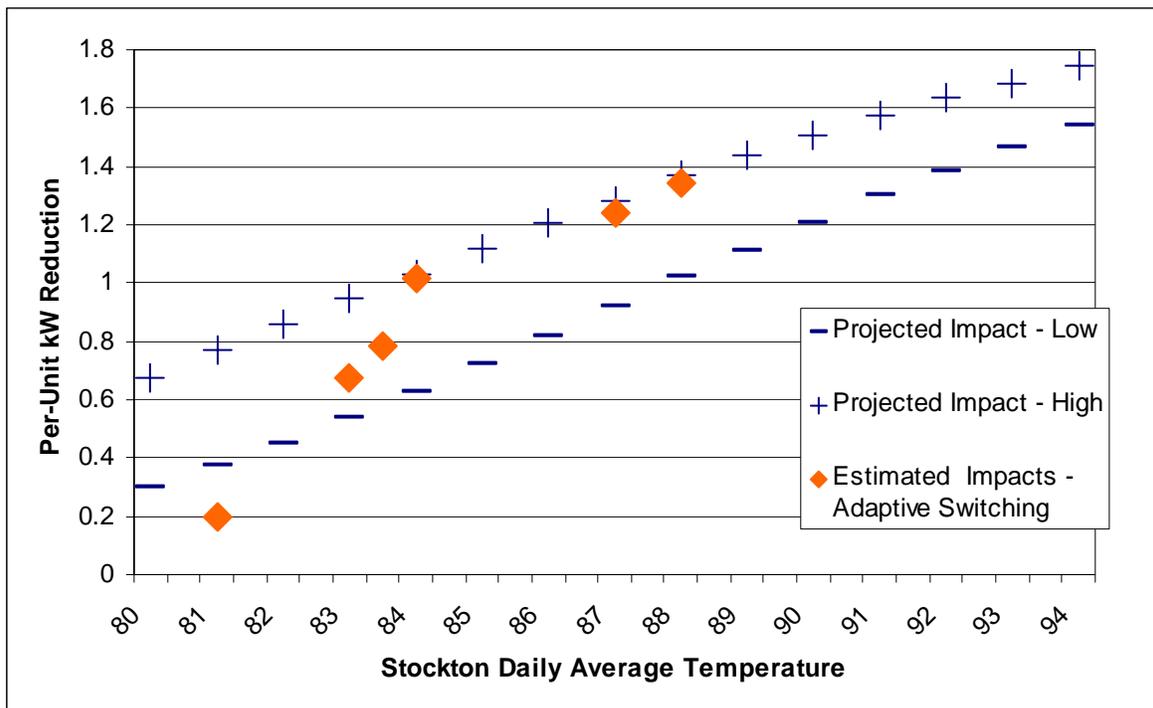
depending on when the control day takes place relative to area weather. More specifically, impacts will reflect individual unit duty cycle on the days that have been chosen for “learning” purposes up to that date.

Our solution to this dilemma is to provide a range of potential impacts. The impacts at the low end of the range reflect the default behavior of the TrueCycle technology; that is, a standard 50 percent control, with no improvements through “learning”. The impacts at the high end of the range reflect effective control at half of the expected duty cycle from our unit-level, hour-specific duty cycle models. This can be thought of as the level of control based on a regression-based learning algorithm using data from the whole summer.

5.5.1.1 Projected Impacts for Switch Participants

Figure 5-8 illustrates the duty cycle-based projections for switch participants. The low projected impacts represent a standard, flat 50 percent control. The high projected impacts represent an adaptive control with an effective control of half the expected duty cycle from our duty cycle models. The observed impacts from the summer 2007 program are shown to illustrate where they fall relative to the projected ranges. Table 5-7 provides the impact projections in tabular form.

Figure 5-8
Impact Projections for Switch Participants with Summer 2007 Estimated Impacts



**Table 5-7
Impact Projections for Switch Participants**

Stockton Daily Average Temperature	Projected Impact - Low	Projected Impact - High
80	0.30	0.68
81	0.37	0.77
82	0.45	0.86
83	0.54	0.95
84	0.63	1.03
85	0.72	1.12
86	0.82	1.20
87	0.92	1.28
88	1.02	1.37
89	1.11	1.44
90	1.21	1.51
91	1.30	1.58
92	1.39	1.63
93	1.46	1.69
94	1.54	1.74

If a real-time learning algorithm is working, it is reasonable to expect an improvement in control over the course of the season. Results earlier in the season will be closer to the low projections as the learning algorithm will have little data from which to diverge from the default 50 percent control. Results later in the cooling season, provided a large number of hot days have taken place, will be closer to the high projections. The results in Figure 5-8 illustrate this trend. The three highest impacts are also the three latest event contained in this figure. They fall substantially closer to the high projections than the other, earlier events.

The high projections produced by the duty cycle models are not designed to reproduce the TrueCycle algorithm. The duty cycle models reflect the unit level data from the whole summer. The TrueCycle algorithm only has learning day data from dates previous to the event. In addition, the SmartAC technology vendor indicates that learning days are integrated into the expected duty cycle with a weight of 1/8. Only after 8 learning days will the default assumptions be completely removed from TrueCycle expected duty cycle. This indicates that the high range projections will tend to over-estimate impacts earlier in the season.

In fact, there is some possibility that both full season duty cycle projections will over-estimate early season impacts. At the same temperatures, there tends to be less cooling early in the season. The duty cycle models have a tendency to project late summer cooling behavior into those early summer periods potentially overestimating expected duty cycle and thus overestimating impacts. The low estimated impact on Figure 5-8 corresponds to July 23rd; it is not an early summer event, but it followed more than two weeks of moderate temperatures.

5.5.2 Thermostat Impact Projections for 2008

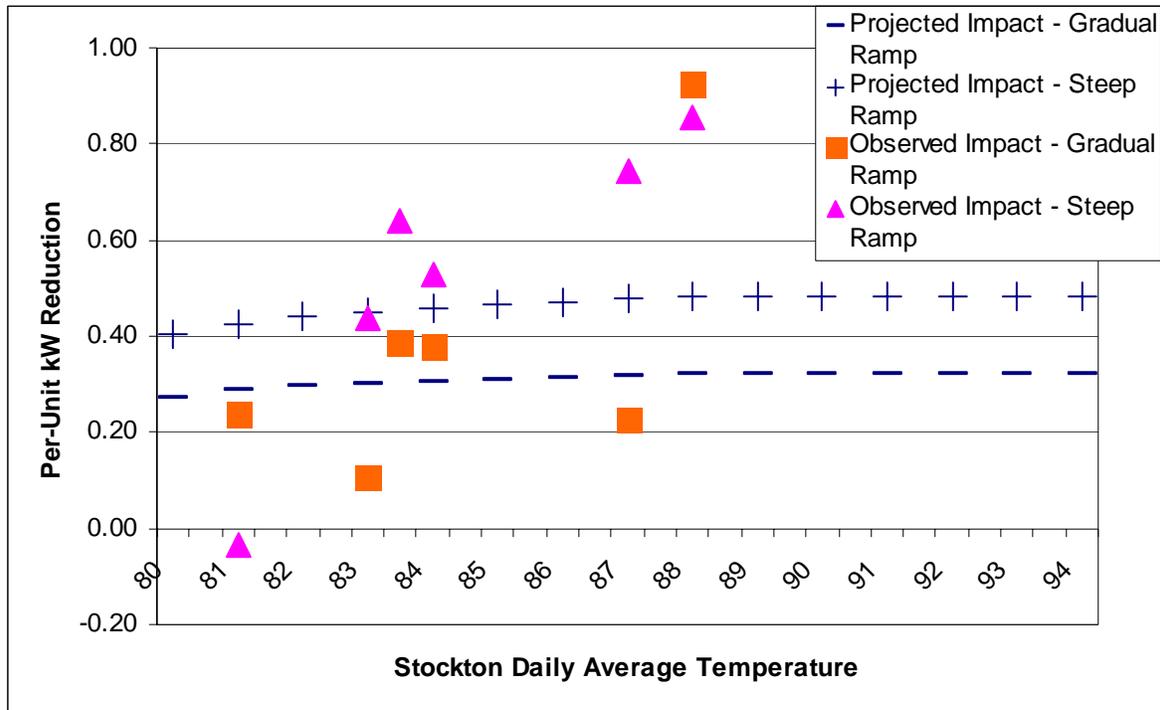
The impact projections for participants with thermostat controls are derived from the kW modeling used to estimate program impacts. The models provide an estimate of average air conditioner load as a function of degree days. The model optimizes the choice of degree day base thereby identifying, on average, the outdoor temperature above which the air conditioner starts to be used. Above that temperature, the model indicates the average load used for cooling for each temperature level.

The impact projections for units with thermostat control uses outdoor temperature differential as a proxy for indoor thermostat setpoint increase. This approach assumes a three degree increase in setpoint is analogous to cooling the house at an outdoor temperature decreased by three degrees. This approach has provided reasonable projections for non-ramped setpoint increases. The innovation needed to accommodate ramping strategies is the use of hour-specific estimates of load differential for different temperature differentials.

The projections addressed two ramping strategies applied to events starting at many different hours. As a solution, for the purpose of projections, KEMA constructed a model event based on the most common start time of events in summer 2007, which was 2 pm. We generated projections for each of the ramping strategies for an event starting at 2 pm.

Figure 5-9 illustrates the kW model-based projections for thermostat participants. The projected impacts of the two ramping strategies are indicated by the crosses and flat lines. The flat lines represent the projected impact for the gradual ramping strategy while the higher crosses represent the steep ramping strategy. Estimated impacts for the two ramping strategies from the summer 2007 program are shown to illustrate where they fall relative to the projected values for each ramping strategy. Table 5-8 provides the impact projections in tabular form.

Figure 5-9
Impact Projections for Thermostat Participants with Summer 2007 Estimated Impacts



**Table 5-8
Impact Projections for Thermostat Participants**

Stockton Daily Average Temperature	Projected Impact - Gradual Ramp	Projected Impact - Steep Ramp
80	0.27	0.40
81	0.29	0.42
82	0.30	0.44
83	0.30	0.45
84	0.30	0.46
85	0.31	0.47
86	0.31	0.47
87	0.32	0.48
88	0.32	0.48
89	0.32	0.48
90	0.32	0.48
91	0.32	0.48
92	0.32	0.48
93	0.32	0.48
94	0.32	0.48

There are two noticeable features of the kW model impact projections compared to the duty cycle model projections:

1. *Unlike the switch unit projections, the projections for both ramping strategies level off above 87°F.* The leveling off of the projected impacts at higher temperatures is an expected feature of setpoint increase-based controls. A setpoint increase, ramped or not, should have the same approximate effect on duty cycle across the range of temperatures at which the unit is running. As expected, load increases with ambient temperature. That increase, however, is gradual at only one percent per one °F outdoor temperature and can easily be masked by the use of daily average temperature as the explanatory variable and the averaging of impacts over many hours. The increase in projected impacts as daily average temperature increases up to 87°F is actually the result of additional units in operation and, thus, being controlled. Above 87°F average daily temperature, all units that ever operated during the summer of 2007 are projected to be operating, and thus impacts level out.
2. *The observed impacts on six hottest days appear to vary widely from the projected levels.* The variability of the observed impact estimates is important to take into consideration. In this sense, Figure 5-9 does not inspire as much confidence as Figure 5-8, above. Only the four event/ ramp strategy combinations with impacts over 0.6 kW (and the one negative event) are statistically different than the projected impacts. It is possible that these more extreme results can be explained by factors outside the capabilities of our models.

The projected impacts reported here clearly underestimated observed impacts on the hottest days of 2007. This points to a need to refine the underlying model that generates the basis for both the estimated and the projected impacts. The data currently available does not support such refinement, but it may be possible when additional data becomes available in the summer of 2008.

5.6 Connected Load

The impact analysis for units with switch control directly estimates unit connected load as a function of hourly temperature. Connected load, the kW draw of a unit when cooling, generally increases linearly with the ambient temperature. Different makes and sizes of units will have unique levels and slopes characterizing their connected load. Section 4.2.2 discusses the derivation of connected load as part of the duty cycle modeling approach.

Table 5-9 presents the population average connected load for three important hourly temperature benchmarks. An hourly temperature of 81°F represents the low end of possible cooling weather. Our analysis found very little cooling activity at daily average temperatures consistent with a maximum temperature of 81°F.

**Table 5-9
Connected Load (kW) by Hourly Temperature**

Hourly Temperature	Connected Load (kW)	Standard Error	Precision at 90 Percent Confidence (90/xx)
81	3.6	0.09	4%
91	3.9	0.09	4%
101	4.1	0.10	4%

A maximum temperature of 91°F is consistent with a daily average temperature of about 79°F. This is the temperature above which the program generated meaningful savings. A maximum temperature of 101°F or more happens only on the most extreme days. Only one day in the summer of 2007 had maximum temperatures over 101°F across the Stockton area. The hottest day of summer 2007, and the event day with the greatest measured impacts, exhibited a maximum temperature of 103°F with a daily average temperature of 88°F. The increase in load across this temperature span (between a daily maximum temperature of 81°F and a daily maximum temperature of 101°F) is 14 percent. This result applies only to units with switch controls that had enough usage to establish a unit-level connected load.

5.7 Snapback

The results provided in this report show that the SmartAC program successfully reduced air conditioning load during the event periods. Positive program impacts indicate that in the absence of the program the air conditioning units would have provided additional cooling. At the end of an event, it is not uncommon for units to compensate for the lost cooling over the previous hours. This can lead to a program-induced increase in load after the event, over what would have happened during those hours had there been no event. This greater than normal load during the post-event period is referred to as snapback.

Table 5-10 reports the program-wide snapback results for all event days. The results are reported in a manner that is consistent with the load reduction results. That is, load reduction is expressed as a positive number while a snapback-related increase in load is expressed as a negative number.

Table 5-10
Estimated Program Snapback results, Average Per-Unit kW per Event(*)

Event Date	Event Day Daily Average Temperature	Event Start	Event Duration	Per Unit Snapback (kW)	Participating Units	Program Wide Snapback (MW)	Confidence Interval Lower Bound (MW)	Precision at 90 Percent Confidence (90/xx)
8/31/2007	88	2 PM	5	0.18	8,843	1.6	2.8	79%
8/30/2007	87	3 PM	4	-0.37	8,809	-3.2	-2.1	35%
8/28/2007	84	3 PM	4	-0.20	8,690	-1.8	-0.6	64%
8/21/2007	84	2 PM	5	-0.26	8,306	-2.2	-1.1	51%
8/22/2007	83	2 PM	5	-0.64	8,391	-5.4	-4.3	20%
7/23/2007	81	12 PM	4	-0.22	6,566	-1.4	-0.6	57%
8/10/2007	78	2 PM	4	-0.08	7,613	-0.6	0.2	> 100%
8/1/2007	77	2 PM	4	-0.68	7,049	-4.8	-3.9	19%
7/27/2007	76	2 PM	4	-0.59	6,825	-4.0	-3.1	23%
8/9/2007	75	2 PM	4	-0.40	7,523	-3.0	-2.2	27%
7/26/2007	75	2 PM	4	-0.49	6,757	-3.3	-2.6	22%
9/26/2007	74	2 PM	5	0.16	10,344	1.7	2.2	31%
9/10/2007	74	4 PM	3	-0.05	9,297	-0.5	0.1	> 100%
7/17/2007	73	12 PM	5	-0.34	6,215	-2.1	-1.4	33%
7/12/2007	71	2:30 PM	3.5	-0.22	5,949	-1.3	-0.8	38%

(*) Snapback period starts 30 minutes after event time and last for 90 minutes.

Snapback was estimated for the 90 minutes following the immediate 30 minutes after the end of the event. The first 30 minutes after the end of the event are removed from the analysis to account for the random start/stop that is implemented by each device.

The estimated snapback results range from an increase of almost 0.68 kW per unit to an apparent continued post-event decrease of load of 0.18 kW. While the load reduction estimates are highly correlated with daily average temperature (the correlation coefficient is 0.92), the snapback results are not (the correlation coefficient is 0.10). This is somewhat counter-intuitive as the post-event cooling load should be driven by temperature as is the load reduction. In fact, on the hottest day and the day with the greatest load reduction, there was no snapback effect. On that day a low level load reduction effect appears to have continued through the snapback period.

These snapback results provide an indication of a potential issue with the weather model estimated load on some of the event days. An over-estimated baseline load will overestimate the load reduction and underestimate the snapback. At the opposite extreme, an underestimated baseline load will underestimate the load reduction but overestimate the snapback. Table 5-11 shows there were four days with daily average temperatures in the mid-70s with almost no recorded load reduction but substantial snapback. There is also a day with an extremely high temperature, and no snapback. These examples indicate that our models underestimate load reduction in the low range of cooling temperatures while they over-estimate load reduction at the highest temperatures. This is consistent with using a linear estimate of a relationship that should have a non-linear shape. There is, in fact, less evidence of this problem when comparing the switch participant results which are modeled non-linearly with a Tobit model of duty cycle.

**Table 5-11
Program Load Reduction vs. Snapback, Per Unit kW**

Event Date	Event Day Daily Average Temperature	Event Start	Event Duration	Program Wide		Thermostat Participants		Switch Participants	
				Impact (MW)	Snapback (MW)	Impact (MW)	Snapback (MW)	Impact (MW)	Snapback (MW)
8/31/2007	88	2 PM	5	1.2	0.2	0.9	0.3	1.3	0.1
8/30/2007	87	3 PM	4	1.0	-0.4	0.5	-0.1	1.2	-0.5
8/28/2007	84	3 PM	4	0.9	-0.2	0.5	0.0	1.0	-0.3
8/21/2007	84	2 PM	5	0.7	-0.3	0.5	-0.3	0.8	-0.3
8/22/2007	83	2 PM	5	0.6	-0.6	0.3	-0.5	0.7	-0.7
7/23/2007	81	12 PM	4	0.2	-0.2	0.1	-0.2	0.2	-0.2
8/10/2007	78	2 PM	4	0.3	-0.1	0.2	-0.3	0.4	0.0
8/1/2007	77	2 PM	4	0.0	-0.7	0.0	-0.9	0.1	-0.6
7/27/2007	76	2 PM	4	0.1	-0.6	0.0	-0.4	0.1	-0.6
8/9/2007	75	2 PM	4	0.1	-0.4	0.2	-0.5	0.1	-0.4
7/26/2007	75	2 PM	4	0.0	-0.5	0.0	-0.6	0.0	-0.5
9/26/2007	74	2 PM	5	0.3	0.2	0.3	0.2	0.3	0.2
9/10/2007	74	4 PM	3	0.1	-0.1	0.1	0.0	0.0	-0.1
7/17/2007	73	12 PM	5	0.0	-0.3	0.0	-0.5	0.0	-0.3
7/12/2007	71	2:30 PM	3.5	0.0	-0.2	-0.1	-0.4	0.0	-0.2

6. Conclusions

The SmartAC program allows participants to choose the control technology with which they will participate in the program. As discussed in section 3, the quite different characteristics of the two technologies are likely to drive participant choice of control technology. As participants self-select into the two control technology groups, those groups may be quite different with respect to participant characteristics that affect AC usage or program performance. In addition, the control strategies applied by the program are difficult to compare ex ante. For both these reasons, comparing results between the two control technologies is not a formal comparison of the effectiveness of the two technologies.

6.1.1 Control Strategy

All reported results indicate a higher level of impact for participants with switch control technology. These results reflect different ex ante control levels. The adaptive cycling was set at 50 percent. The most common cycling strategies range from 33 to 50 percent. The level chosen for use by SmartAC was at the higher end of similar DLC programs. In addition, the adaptive nature of the TrueCycle technology only improves the efficacy of the switch control relative to the standard 50 percent cycling regime.

The thermostat controls used ramping strategies. Both strategies increased in single degree increments and topped out at a maximum of a three degree set point increase. The “steep” ramping strategy reached the three degree set point increase in the third hour. The “gradual” strategy reached the third degree in the fifth hour, when the control event was about to end, or not at all in situations where the event lasted four hours or less. Ramping strategies are still relatively new. They are used either to improve participant comfort or smooth the load impacts over the event, or both. The hourly impact results for the combined ramping strategies indicate that thermostat impacts were relatively lower than switch impacts during early event hours but were closer to switch results by the end of the events. This could indicate too gradual a ramp. Regardless, a three degree set point increase, steep or not, is only a moderate setpoint increase.

The thermostat controls clearly provided lower load impacts than the switch controls. These results reflect the relative aggressiveness of the control strategies used for the two technologies. In theory, the thermostats could be controlled so as to produce impacts commensurate with the switch results reported here. PG&E SmartAC applied switch control levels that are relatively aggressive compared to other similar programs. Conversely, it also applied relatively moderate thermostat control levels.

6.2 Participant Satisfaction

One possible point of comparison for the two technologies is load impact relative to participant satisfaction. As part of the Process Evaluation conducted for the SmartAC’s first year¹¹, participants were asked about their awareness of and level of comfort during the events.

Overall, only 40% of respondents were aware PG&E had activated their AC control technology. The overall result breaks out into statistically significantly different results when considered by control

¹¹ Final Report: Process Evaluation of 2007 PG&E Smart AC Program. Study ID PGE0262.01. March 31, 2008

technology. Forty-four percent of those with switch controls reported being aware of the events while only 28 percent of those participants with thermostat controls reported being aware of the events.

Those respondents who were aware of the events were asked about comfort level. Overall, 25 percent said they were “somewhat” or “very” uncomfortable during the control events. This overall result also breaks out into statistically significant different results when considered by control technology. Twenty percent of those with switch controls reported being “somewhat” or “very” uncomfortable while 49 percent of those participants with thermostat controls reported this level of comfort.

The best way to combine these potentially confounding results is to isolate those participants expressing dissatisfaction with the level of comfort. Combining results from the two questions above the overall percentage of participants expressing discomfort from thermostat and switch groups was 14 and 9 percent, respectively. This difference is not statistically significant. Thus, despite having a higher level of load impact, switch participants did not indicate higher levels of discomfort.

Comparing participant satisfaction across the groups can only be done with the caveat that these two groups are self-selected and may have underlying differences that drive the differences. While the survey results do indicate clear differences in awareness and comfort levels, the overall percentage of respondents both aware of the events and experiencing discomfort were not statistically different.

6.3 Switch Summary

Direct load control switches made up approximately 70 percent of the SmartAC population at the time of the 2007 system peak. The SmartAC evaluation results indicate substantial savings for the participants with switches on the hot event days. On the three event days surrounding the day of the system peak switch participants registered a minimum of 0.9 kW load reduction during the peak hours of the day. Results indicate that the adaptive switching technology performed better than if the standard 50 percent switching had been employed.

There is some suggestive evidence that, despite the adaptive switch, the switch technology is still sensitive to AC unit oversizing. There is also tentative evidence that the adaptive switch achieved higher impacts with no increase in expressed discomfort. Further resolution of these questions is left open for future research.

6.4 Thermostat Summary

Thermostat demand control devices made up approximately 30 percent of the SmartAC population at the time of the 2007 system peak. Thermostat control results indicate demand reduction substantially below that realized by switch participants. On the three event days surrounding the day of the system peak thermostat participants registered load reduction ranging from zero to over 1.3 kW during the peak hours of the day. We believe that the difference in impacts is explained by a combination of the choice of ex ante control level employed by the program and potential population differences resulting from self-selection. These results should not be construed as necessarily reflecting the relative efficacy of thermostat controls versus switches.

Results comparing the two different ramping strategies provide some evidence of increased impact from the steeper strategy. There is suggestive evidence that the thermostat control is less sensitive to

oversizing. On the other hand, there is also evidence that despite lower levels of impact, thermostat participants had similar levels of discomfort as switch participants. Once again, further resolution of these questions is left open for future research.

7. Appendix A

7.1 Tobit Model of Duty Cycles

Using duty cycles derived from the one-minute interval data, we developed models to describe duty cycle as a function of weather. These models necessarily excluded all duty cycle observations on weekends and holidays, as well as all event days.

The underlying relationship between natural duty cycle and weather is a linear one, but it is constrained by practical duty cycle limits of one and zero. A simple linear regression thus does not apply. To identify the underlying linear relationship while still accounting for the practical constraints to duty cycle, we employed a Tobit analysis. A Tobit analysis explicitly allows for the dependent variable in a linear relationship to be constrained, whether at lower, upper, or both ends of some range. The products of the Tobit analysis include a linear expression of duty cycle and scale factor by which the probability associated with a particular duty cycle can be determined.

The Tobit analysis models a theoretical “unconstrained” duty cycle as a function of the weather variable. The unconstrained duty cycle is not limited by the practical limits to a duty cycle of zero and one, but instead can be negative or greater than one. The observed duty cycle is the result of bringing the unconstrained values into the physically possible bounds of 0 to 1, as in Equation 11.

Equation 11

$$DC_{ip} = \begin{cases} 0, & DC_{ip}^* \leq 0 \\ 1, & DC_{ip}^* \geq 1 \\ DC_{ip}^*, & 0 < DC_{ip}^* < 1 \end{cases}$$

where

$$DC_{ip}^* = \text{unconstrained duty cycle during period } p \text{ for CAC } i.$$

We fit multiple models for each unit to account for changes in cooling behavior over the course of the day. For this evaluation there was sufficient data to estimate hourly models. For each unit, we estimated hourly duty cycle models for the hours between noon and 10 PM. All the equations in this section are expressed with p for time period. For this evaluation, the period was hourly.

The models were initially fit with the four weather variables described in section 4.2.3, to identify the variable that would give the best fit. The model of unconstrained duty cycle as a function of a temperature variable is given by Equation 12. For this equation and those that follow, W_d could be any daily temperature or *THI* variable.

Equation 12

$$DC_{ip}^* = \alpha_{ip} + \beta_{ip} \times W_d + \varepsilon_{ip}$$

where

- DC_{ip}^* = unconstrained duty cycle during hour h of period p for CAC i ,
- α_{ip}, β_{ip} = Coefficients for period p and CAC i ,
- W_d = daily temperature variable for day d , and
- ε_{ip} = residuals independently and normally distributed in period p for CAC i , with mean equal to zero and variance equal to $\hat{\sigma}_{ip}^2$.

We developed estimates of the coefficients in Equation 13, \hat{a}_{ip} and \hat{b}_{ip} , and estimates of the variance, $\hat{\sigma}_{ip}^2$, using standard Tobit estimation procedures.

Equation 13

$$\Phi(x) = \int_{-\infty}^x \phi(u) du = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du$$

The duty cycle models allow unbiased estimates of natural duty cycle to be made for different weather variable levels. The calculation uses Equation 14.

Equation 14

$$\begin{aligned} E(DC_{ipw}) &= E(DC_{ipw} | 0 < DC_{ipw} < 1) \times \Pr(0 < DC_{ipw} < 1) \\ &\quad + E(DC_{ipw} | DC_{ipw} = 1) \times \Pr(DC_{ipw} = 1) \\ &= \hat{DC}_{ipw}^* \times \{\Phi(U) - \Phi(L)\} + \hat{\sigma}_{ip} \times \{\phi(L) - \phi(U)\} + 1 - \Phi(U) \end{aligned}$$

where

- $E(DC_{ipw})$ = estimated natural duty cycle in period p at temperature w for CAC i ,
- \hat{DC}_{ipw}^* = estimated unconstrained duty cycle in period p at temperature w for CAC i ,
- Φ indicates the standard normal cumulative distribution function, and
- ϕ indicates the standard normal probability density function.

The expressions L and U in the preceding equation refer to the lower and upper practical constraints on duty cycle, here being zero and one respectively. They are defined by the following two equations:

Equation 15

$$L = \frac{0 - \hat{DC}_{ip}}{\hat{\sigma}_{ip}}$$

Equation 16

$$U = \frac{1 - \hat{DC}_{ip}}{\hat{\sigma}_{ip}}$$

7.2 Individual Demand Impacts by Duty Cycle Approach

The duty cycle modeling approach also provides estimated duty cycle reduction as a function of actual control level. This is easiest to explain given a constant control level. A constant 50 percent control level has been widely used by switch programs. The remainder of this discussion focuses on the estimated duty cycle reduction as a result of a 50 percent control level, however, the same basic technique can be employed for any level of control.

The Tobit analysis estimates duty cycle within certain constraints given an underlying linear relationship with weather. The natural duty cycle estimate above was constrained to stay between a duty cycle of zero and 100. The model takes into consideration the probability that the expected underlying linear duty cycle would fall outside the constraints. The same process takes place when estimating duty cycle given a control level of 50 percent. The control level of 50 percent provides the new minimum constraint.

The coefficients from the Tobit model fit $(\hat{a}_{ip}, \hat{b}_{ip}, \hat{\sigma}_{ip}^2)$ were used to calculate for each CAC in different time periods p and at different temperatures W :

- the probability that $DC_{ipw} \geq 0.5$, and
- the expected value of DC_{ipw} , given that $DC_{ipw} \geq 0.5$.

The probability and the expected value of the duty cycle, given that the duty cycle is greater than 0.5, were calculated using the following two equations:

Equation 17

$$\Pr(DC_{ipw} \geq 0.5) = 1 - \Phi(L)$$

Equation 18

$$\begin{aligned} E(DC_{ipw} | DC_{ipw} \geq 0.5) &= E(DC_{ipw} | 0.5 \leq DC_{ipw} < 1) \times \Pr(0.5 \leq DC_{ipw} < 1 | DC_{ipw} \geq 0.5) \\ &\quad + E(DC_{ipw} | DC_{ipw} = 1) \times \Pr(DC_{ipw} = 1 | DC_{ipw} \geq 0.5) \\ &= \left(\frac{1}{1 - \Phi(L)} \right) \times \left[\hat{DC}_{ipw}^* \times \{\Phi(U) - \Phi(L)\} + \sigma_{ip} \times \{\phi(L) - \phi(U)\} + \{1 - \Phi(U)\} \right] \end{aligned}$$

where

$$E(DC_{ipw} | DC_{ipw} \geq 0.5) = \text{expected duty cycle in period } p \text{ at temperature variable } w \text{ for CAC } i \text{ given that the duty cycle is at least } 0.5.$$

The expressions L and U in the two preceding equations refer to the lower and upper practical constraints being placed on duty cycle. The equation for U is the same as used Equation 16 where the upper duty cycle constraint is one. The equation for L here, however, uses a lower duty cycle constraint of 0.5 rather than zero as was used in Equation 15 with reference to an unconditional estimate of duty cycle. The expression for L for an estimate of duty cycle conditional on being greater than 0.5 is given by Equation 19.

Equation 19

$$L = \frac{0.5 - \hat{DC}_{ip}}{\hat{\sigma}_{ip}}$$

The expected duty cycle given a minimum duty cycle of 50 percent represents the potential duty cycle reduction if limited to a maximum duty cycle of 50 percent. The expected duty cycle given a minimum duty cycle of 50 percent will remain above 50 percent despite a natural duty cycle below 50 percent because there remains some probability that natural duty cycle is, in fact, greater than 50 percent. The estimate of the natural duty cycle is a mean estimate with a surrounding probability distribution. Even when the mean estimate of the natural duty cycle as a given temperature is less than 50 percent, there still remains part of the probability distribution above the 50 percent cut-off.

Finally, the projected load reduction for 50 percent control is calculated by

Equation 20

$$\bar{S}_{ipw} = \hat{C}L_{iw} \times (\hat{DC}_{ipw, > 0.5} - 0.5)$$

where

$$\bar{S}_{ipw} = \text{Load reduction in period } p \text{ at temperature } w \text{ for CAC } i, \text{ and}$$

$\hat{DC}_{ipw,>0.5}$ = expected natural duty cycle given that it is greater than 0.5 in period p at temperature w for CAC i .

Any level of control, including an adaptive control that is a function of the natural duty cycle can be calculated by inserting the related lower duty cycle constraint for the 0.5 used in these equations.