Using Residential AC Load Control in Grid Operations: PG&E's Ancillary Service Pilot

Michael Sullivan, Josh Bode, Bashar Kellow, Sarah Woehleke, and Joseph Eto, *Member, IEEE*

***Abstract*—This paper summarizes the results from a PG&E pilot designed to assess the ability to use air conditioner direct load control to provide ancillary services. The study included nearly 2,000 residential households, with control devices that were instructed to cause an immediate and complete shutdown of the air conditioner compressor 71 times over a two and half month period. It summarizes the start and total ramp time of AC load response, the magnitude of the response, the effect of the curtailments on customers comfort and satisfaction, and the approach to providing near real time visibility of the air conditioner electricity demand. [1]**

***Index Terms*—Ancillary services, air conditioner, direct load control, ramp speed, start time, reserves, air conditioner, load management.**

I. Introduction

Aggregate electricity consumption and production must match in real time in order to maintain the stability and reliability of the electric grid. Because generation and transmission facilities can be unexpectedly forced out of operation, system operators maintain controllable reserves to meet short term fluctuations in the supply-demand balance due to unforeseen generation outages, transmission outages, forecast error, and volatility in intermittent generation.

There are three types of reserves – frequency regulation, 10-minute spinning and non-spinning reserve and 30-minute operating reserves. Regulation reserves are synchronized with the electric grid and respond continuously over very short-time scales (i.e., seconds) to balance voltage and maintain power quality. Historically, contingency reserves have been supplied by generating machines that are synchronized with the load on the grid but have additional generating capacity above their current dispatch point (spinning reserves) or by generators with fast start capabilities (non-spinning reserves). The defining characteristics of contingency reserves are the speed of their response and ramping speed. Figure 1 describes the role of reserves in ensuring system reliability.

This paper first submitted on June 8, 2012.

M. Sullivan, J. Bode, and S. Woehleke are with The FSC Group, San Francisco, CA 94104.
J. Eto is with Lawrence Berkeley National Laboratory, Berkeley, 94720.
B. Kellow is with Pacific Gas and Electric, San Francisco, CA 94104.

Because contingency reserves are used to back-up the system, they are operated infrequently and, typically, for less than 10 minutes at a time. In 2009, Kueck and Kirby reviewed the frequency and duration of contingency reserve operations by the California, New England, and New York Independent System Operators (ISO). They found that deployments over 30 minutes were very rarely needed and that contingency reserve averaged roughly 10 minutes in each of the ISO's [2].

Figure 1: The role of spinning and non-spinning reserves

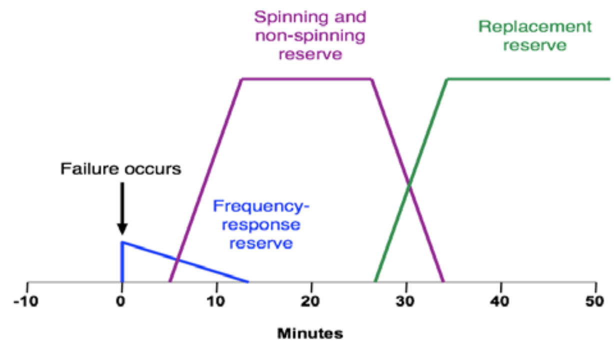

Some types of DR are capable of providing these reserve services at lower cost and with better performance than conventional generating facilities. This is accomplished by rapidly decreasing electricity demand rather than rapidly increasing generation.

Interest in using air conditioner load response for system operations has increased in the past decade due to several reasons, including:

- The need for fast-ramping resources in electricity systems has increased.
- There are many pre-existing air conditioner load control programs with substantial unused capacity that, to date, have not been incorporated into grid operations or ancillary service markets. Based on the Federal Energy Regulatory Commission 2010 DR survey [3], there are over 200 such residential programs with over 4.8 million households enrolled.
- There are no fuel costs and no emissions associated with the operation of these resources.
- Recent technological innovations enable aggregation of many small scale loads.

- Recently developed optimization algorithms allow for more precise load control.
- The technology to provide near real time visibility of loads is available.
- Many load control devices now include over and under frequency relays, providing an automated fail-safe mechanism that is synchronized with the electricity system.

Interest on using loads rather than generators to provide contingency reserves and ramp speed date back over a decade. Several demonstrations and pilots have shown that air conditioner loads can be used during the summer cooling season to improve the stability and reliability of the electricity grid by providing ancillary services.

In 2001, Eto presented a research agenda for improving reliability by using customer loads to provide ancillary services. [4] In 2002, Kueck and Kirby began testing the ability to systematically control air conditioners to provide contingency reserves with 24 units at a hotel. [5] Subsequently, they tested the use of water pumps to provide fast response resources. [6] The first large scale test of using air conditioner direct load control to provide contingency reserves occurred in Southern California Edison (SCE) in 2006. [7] In total, the compressors from approximately 270 air conditioner units were curtailed in full for 5 minutes for a total of 37 times. SCE expanded the pilot in 2008 to include 1,200 households and 50 load curtailments; the vast majority of which lasted no more than 6 minutes each. [8] [9]

This paper presents the results from a pilot implemented by Pacific Gas and Electric (PG&E) in 2009 designed to assess the ability to use air conditioner direct load control to provide ancillary services. The pilot included nearly 2,000 households concentrated in 4 feeders, each of which had their AC units curtailed for between 68 and 71 times each, for 15 minutes at a time. The main objectives of the pilot were to:

- Simulate the provision of spinning and non-spinning reserve with air conditioner load control.
- Assess how much time elapsed before air conditioning units started reducing loads and how long it took to ramp up to full capacity.
- Test out the ability to provide near real time visibility of load.
- Determine if the demand reductions observed with the sample were also observable with feeder data.
- Assess the magnitude and variation in the demand reductions.
- Determine if short duration air conditioner curtailments affected customer comfort and / or led to participation fatigue.

The remainder of this paper presents the pilot design and results. Section II summarizes the pilot design and the testing protocols. Section III documents start time and ramping speed of air conditioner load curtailments. Section IV summarizes the demand reductions observed, including the ability to provide near real time visibility. Section V documents the impact of repeated short duration curtailments on customer comfort.

II. Pilot Design and Testing Protocols

Four distribution feeder circuits were chosen for study – two in a relatively hot climate (Fresno) and two in warmer Bay Area suburbs (Antioch and Fairfield). The difference in climate is important in that past studies have indicated air conditioner demand and load control performance varies dramatically with ambient temperature.

The feeders were selected for the study because they met key criteria to meet the research objectives, specifically:

- One-minute feeder measurements (KW, Amps, MVAR and temperatures) could be accessed through a secure data port maintained by PG&E.
- The circuit did not contain large commercial and industrial loads.
- A sufficient number of participants on PG&E's direct load control program, SmartAC, already existed on the feeder so that recruiting goals could be met quickly (within one month).

The number of direct load control participants was increased on each of the selected feeders to ensure that approximately 500 SmartAC customers were present. Two technologies with the capability of shutting down the air conditioner compressor remotely were used: load control switches and programmable communicating thermostats. For each feeder, the goal was to include 400 AC units with direct load control switches and about 100 AC units with programmable communicating thermostats. Table 1 summarizes the number of participants in each feeder for the pilot.

Table 1: Pilot AC load control participants

Feeder	Residential Premises	Control Switches	T-stats	Total	Feeder Penetration
Fresno 1	3,250	366	86	452	13.92%
Fresno 2	2,750	420	105	525	19.06%
Antioch	2,900	413	102	515	17.73%
Fairfield	4,740	406	96	502	10.59%
TOTAL	13,646	1,605	389	1,994	14.61%

Figure 2: Screenshot of the output from the telemetry system

The usage was recorded for each of the air conditioner units through the course of the study. For 129 randomly selected units, data was recorded at one-minute intervals using technology that allowed for near real-time transmittal of the air conditioner end use data. Data for the remaining units was recorded at five-minute intervals using standard end-use data recorders without the capability of transmitting data on a near real-time basis.

PG&E's direct load control system was programmed to cause a complete shutdown of all air conditioner compressors for the participants on the feeders under study. Simulated ancillary service operations were conducted on all four feeders simultaneously, twice each week day, at varying hours between 12:00 pm and 7:00 pm. For each test operation, the air conditioner compressors were instructed to shut down for a 15-minute period; control of the units was released over a random 2-minute interval after the test event period concluded. After the testing period, instructions to shut down air conditioner compressors were re-sent five minutes after the initial transmittal in case individual units did not receive the initial curtailment instructions.

During the study period, 71 test control operations were conducted, producing observation of load control over a wide range of times and temperatures. The first 10 test operations were used to refine the dispatch procedures and ensure the devices were correctly programmed. Of the 71 operations, 68 included all 4 feeders and all air conditioner units in the program. For each test operation, the meters with real-time transmittal capability reported the measurements for the 15 minutes before, during, and after each test operation. The data was integrated with feeder load data and displayed live via the internet. Both the California System Operator and PG&E's operations department were able to view air conditioner and feeder demand levels before, during, and after each test operation.

Figure 2 displays a screen shot of the output from the telemetry system for 1 of the 71 test operations. The screen updated every minute based on the measurements taken in the prior minute. The system displayed a graph of the load measurements from the feeder (top left corner); a graph of the load measurements from the sample (bottom left corner); the sample load impacts scaled for the number of participating air conditioner units (top right corner); and useful statistics describing the load response (e.g., average load impact per control unit, percent of appliances in operation, etc.). Users could at any time select which feeder to view or jointly view the loads for all feeders and all sampled air conditioner units on a summary tab. With the system, operators were able to determine how long it took for control to take effect, how long it took for loads to come under full control, and the overall magnitude of load reduction that was achieved.

In the example display screenshot, the load control operation is observable on both the feeder measurements (upper left quadrant) and in the sample observations (lower left quadrant). The impacts could be clearly observed on feeder loads on hotter days when most air conditioner units were operating. This was not true for cooler days when fewer units were in operation. It is also important to recall that the feeders were selected precisely because of the high penetration of air conditioner load control and the availability of one-minute feeder data. Most feeders have a lower penetration of air comparison load control, in which case AC operations are harder to visually detect with feeder data, although they were clearly visible in the sample.

The large number of test operations provided the ability to estimate the load impacts that could be obtained from air conditioner direct load control at different times of day and under different temperature conditions. It also allowed for a detailed assessment of when and how frequently curtailments could be directly observed on the feeder loads. In addition, it provided data from the nearly 2,000 customers on whether or not repeated short AC load control operations affected their comfort levels.

III. Start Time and Ramp Speed

The start time, ramp speed, and ability to be synchronized with the electric grid determine the usefulness of resources for grid operations. Delays can occur at the following stages of the communication:

- Creation of the control message sent to devices;
- Connecting to the outgoing modem used to transmit the signal to paging companies;
- Acknowledgement of the load control signal by the paging systems;
- Transmittal of the signals from the paging network to the control devices; and
- Receipt of the load control signal by the devices.

As noted earlier, after the first few test operations, it was clear that not all devices responded to the test operation signal. In an effort to eliminate this problem, the load control system was programmed to resend the control signal after five minutes to increase the likelihood that devices received the intended signals.

The amount of delay associated with the first three reasons was directly measured using communication logs maintained by the AC load control operation system. However, it was not possible to directly measure the time required for the paging company to transmit the signals to the control devices because communications logs for the paging system were not available to PG&E.

To complicate matters, air conditioner units constantly turn on or off to regulate the temperature inside homes. Controlling air conditioners has two effects: the compressors that are turned on shut down earlier than they would otherwise; and units that are off are prevented from turning on. The amount of time an air conditioning unit is in operation over the course of an hour varies depending on weather conditions. Viewed at one-second time periods, the AC compressor and fan is either on or off. The air conditioner load can take on three states – full load, fan load, and no load. As a result, it is necessary to distinguish the normal patterns of air conditioning units turning off from instances where the AC compressors shut down due to load control operations. To do so, we relied on the sample of air conditioner units with one-minute data and a statistical technique known as survival analysis (also known as time-to-event and time-to-failure analysis). The survival analysis was designed to answer two questions: how long after the start of the test event does it take for AC units to respond to the control signals; and when is the full impact of the load control reached?

To distinguish normal patterns of air conditioning units shutting down from load control operations, we used the 15 minutes prior to each test operation as a control period. The weather conditions, occupancy patterns, and participant characteristics during the 15 minutes *immediately prior* to the test operation are similar to those during the *actual* test operation. In each case, we took a snapshot of the number of units on at the start of the period and determined the share of them that remained as time elapsed. By comparing the rates at which AC units shut down for these two periods, it was possible to identify the time (in seconds) elapsed before the load control operations induced AC units to shut down.

An example of this process is presented graphically in Figure 3, which displays the first 200 seconds of the 15 minutes before and after a control operation test. During the pre-event period, the share of units that remained on decreased gradually – after 60 seconds, roughly 95% of air conditioning units remained on. In contrast, share of units that remain on does not decrease gradually during the test event period. Rather, there is a noticeable drop in the share of air conditioner units that remain on between the 45th and 60th seconds after operations were started.

Figure 3: Example of start time and total ramp time analysis

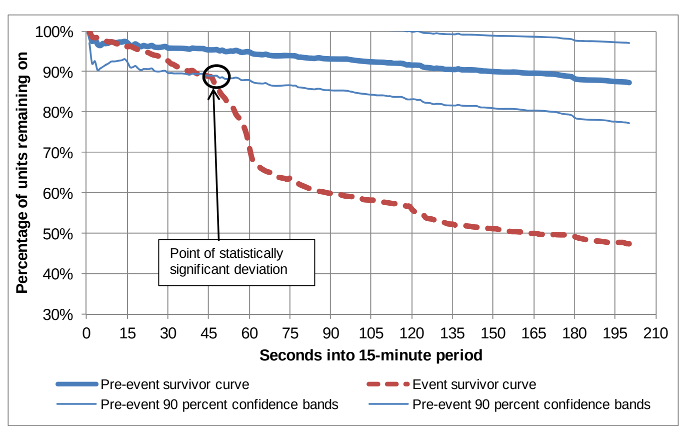

While the share of air conditioner units remaining on starts to deviate for the event and control period about 15 seconds into the event, the difference is not statistically significant until the 47th second. In the analysis, the load control start time was recorded as the time period when differences in the share of air conditioner units remaining on became statistically significant. This was calculated separately for each event. The approach produces a conservative estimate of time-to-start (it over estimates it) because it relies on the amount of time until the difference is statistically significant.

The analysis of the time elapsed between dispatch instructions and the resource start time excluded test operations that had less than 30 air conditioner units in operation at the start of the period. Simply put, the statistical technique used is not reliable when the samples are too small. In addition, it excluded the initial five days when final operation protocols were still being determined. As a result, start times and total ramp times were only recorded 49 out of the 71 test operations.

Figure 4 summarizes the amount of time until air conditioner units noticeably began to shed load for each of the events (start time). The figure separately plots the time until the operation signal was sent out from the AC load control system and the time elapsed between transmittal from the system and receipt of the signal by the load control devices.

The median time to response was 60 seconds. By 120 seconds (2 minutes), responses had occurred in 92% of the tests; responses had occurred in 95% of the tests within 126 seconds. The average time to response was 69.4 ± 8.9 seconds.

Figure 4: Time elapsed until start of AC operations by source of delay

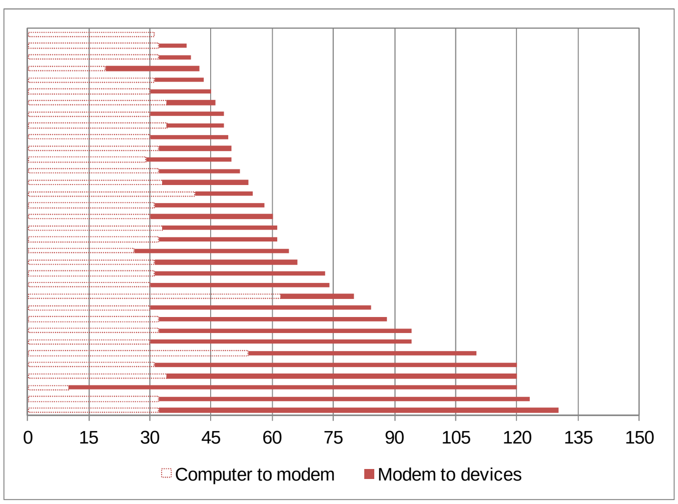

While the start time matters, it is just as important to understand how long it took to bring all of the devices that did respond under control. In other words, how long does it take to ramp up air conditioner load control to full capability? Does the speed of the response vary substantially or is it consistent?

To determine the total ramp time, we relied on the same 129 air conditioner units with 1-minute interval data, except that the focus was on how the rate at which resources came online. Figure 5 illustrates the time until load reductions reached the full reduction capability. Because air conditioner demand levels differ for each event due to weather and time of day, the loads were normalized to the demands observed in the five minutes immediately preceding each event. Put differently, the graph shows the air conditioner electricity demand during the event as a percentage of the air conditioner electricity demand immediately prior to the event. The graph excludes the same events that were excluded from the analysis of load control start time, namely: mild days when demand reductions were not observed due to the lack of air conditioner use, and the week when operating procedures were being finalized.

During test operations, the curtailment instructions caused most air conditioner units to shut off within the first three minutes of dispatch instructions. On average, by the second minute, 65% of the demand reduction had been attained. By the third minute after the dispatch, on average, 80% of the resources were available. Thereafter, the ramp rate for the remaining air conditioner load reduction resources slowed down. By the tenth minute, on average, 98% of the attainable reductions had been achieved. The demand reduction from a centrally dispatched system was not immediate for all events. In few instances, the demand reductions took up to seven minutes. These were generally instances where delays were experienced in the communication time from the load control operating system to the devices.

Figure 5: Air Conditioner Ramp Speed by Event

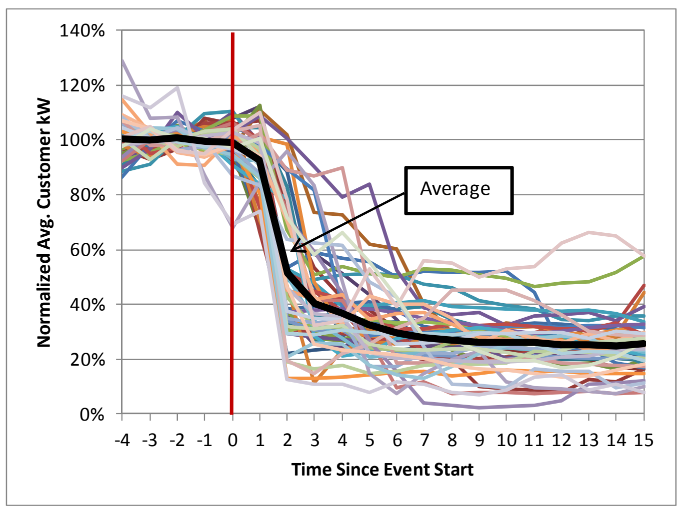

The pilot did not test the over and under frequency relays built into PG&E's switch devices. This feature provides an automated fail-safe mechanism that is synchronized with the electricity system. When the system frequency (or voltage) drops below a specified threshold (which is adjustable), it automatically shuts down the air conditioner compressor. This failsafe mechanism does not rely on centrally dispatched instructions. In theory, it should be able to reduce demand nearly instantaneously, if needed.

IV. Demand Reductions

Because reserve operations are of a very short duration, it is possible to completely shut off the air conditioner compressor.

This produces larger demand reductions than traditional load control operations, which typically only allows for partial air conditioner control and cuts back the operation of the air conditioner compressor by about half.

Air conditioner loads vary substantially with ambient temperature and time of day and directly affects the amount of reserves that can be provided. Figure 6 illustrates the variation in AC use for three days in September with daily maximum temperatures of 86°F, 91°F, and 100°F. The electric load for the 91°F day is about double the load on the 86°F day and the load for the 100°F day is almost three and half times larger. In addition, on all three days, the air conditioner demand varied substantially by hour. The downward spikes in the graph reflect curtailment operations. The graph also shows the estimated air conditioner load had the units not been curtailed, known as the reference load or baseline. We discuss the impact estimation in more detail below.

Clearly, the amount of contingency reserves that residential air conditioner units can provide are dependent on the weather conditions and the time of day. While air conditioning demand is a variable resource, it is highly predictable and dispatchable. Importantly, its availability increases precisely when system loads are highest and resources are most needed.

Figure 6: Example variation of AC loads by time of day and temperature

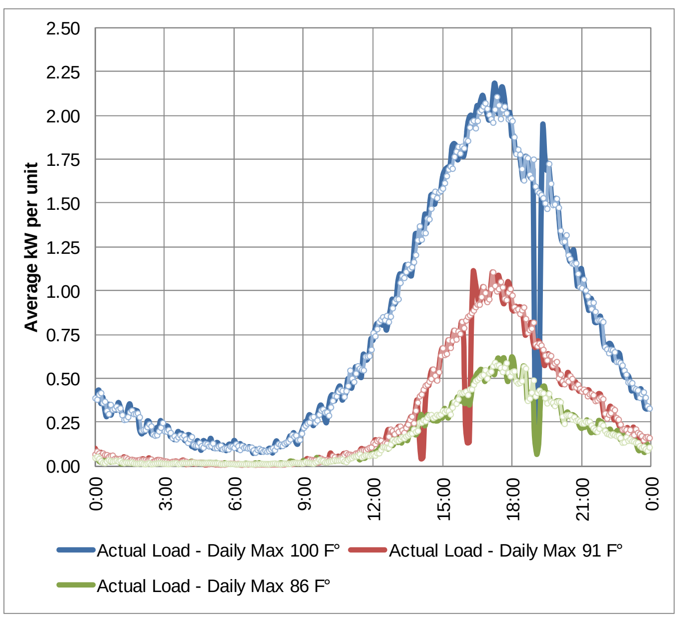

Just like air conditioner load, impacts vary substantially by time of day and hour. Figure 7 shows the distribution of per device demand reductions observed across all test operations. The events were called twice each weekday over the test period at random start times. The average load reduction per device across all events was 0.65 kW per device. However, for many periods, the reductions per air conditioning unit were small and near zero. On the other hand, the demand reductions exceeded 0.5 kW per air conditioner unit for 60% of the test periods and 1.0 kW per unit for 30% of the test periods. Aggregated over 150,000 units, it can be a substantial resource that can be quickly started and ramped up.

Figure 7: Distribution of test operation demand reductions

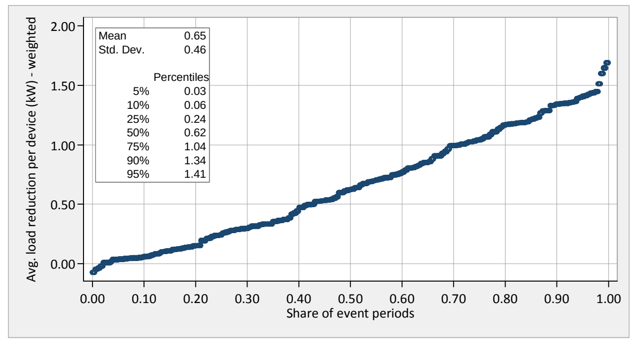

Because spinning reserves are short in duration, the air conditioner demand absent curtailment operations can be estimated based on the usage patterns observed in periods immediately before and shortly after each operation. The customer is not directly notified of operation, so they are unlikely to change their behavior and influence the baseline or the magnitude of demand reductions. This approach was implemented using aggregated air conditioner load from the sampled participants in each feeder. The estimates were produced using regression because it automated the process and produced standard errors that corrected for auto-correlation. The model explained most of the variation in air conditioner loads, as can be seen in Figure 4, presented earlier, which compares actual measured AC loads to the counterfactual predicted by the model. The predicted reference loads mirror the actual loads for all periods, except those affected by the curtailment test operations.

Table 2 presents the estimated demand reductions by temperature range for each of the four feeders in the pilot. Air conditioner loads in the 95-100°F temperature range are typically more than 50% higher than loads in the 90-95°F range. As a result, the load reduction potential from air conditioner control varies substantially, with more resources becoming available on the hotter days that typically drive system peaking conditions for most utilities.

The results in Table 2 also reflect differences in impacts across feeders. Several factors explain those differences. The air conditioner loads reflect not only temperature conditions, but the size of homes (which determines the size of the units), age of homes, household characteristics, and operation patterns.

In the milder regions (relative to Fresno) of Antioch and Fairfield, participants tended to use air conditioners at lower temperatures while Fresno participants typically increased air conditioning consumption when temperatures exceed 90°F. In addition, the reductions as percentage of the baseline also varied. In specific, approximately 80% of the air conditioning electric demand was reduced during operations in the Antioch and Fairfield feeders, with some variation across events. However, on average, less than 60% of the air conditioner electric demand was reduced during operations in the two Fresno feeders. This pattern was identified via the real time monitoring system shortly after operations started and was investigated by measuring the strength of the paging signals in each area. The difference in the demand reductions largely was due to weaker paging network in the areas where the Fresno feeders were located, even though curtailment signals were sent over two paging systems to ensure robust coverage.

The variation in the communication network coverage also has implications for the type of devices utilized. Load control switch devices are typically situated outside homes and can more easily receive and respond to weaker paging signals. On the other hand, thermostats are typically inside homes and are less likely to receive and respond to weaker paging signals. Both switch devices and thermostats were operated on shed mode and, in theory, should have produced similar demand reductions if they received the curtailment signal. However, the percent of air conditioner load curtailed varied by device types within each feeder. In the Fresno I feeder, on average, switch devices reduced loads by 71%, while thermostats only reduced them by 39%. In the Fresno II feeder, the difference was even larger. Switch devices reduced air conditioner loads by 77%, on average, while thermostats only curtailed 28% of the load. The differences between thermostat and switch device demand reductions were smaller for two feeders with stronger paging network signals. In the Antioch feeder, on average, switch devices reduced loads by 82% and thermostats delivered 69% reductions. In Fairfield, switch devices and thermostats reduced loads by 84% and 67%, respectively.

V. Effect of Load Curtailments on Customer Comfort

A key concern of using end use loads for ancillary services is customer reaction to repeated, short-duration operations. The PG&E pilot was one of the first tests that allowed large scale measurement of how customers reacted to frequent and short curtailment operations. No other tests to date have called as many operations as frequently and for as many customers as was done for this pilot.

The effect of the test operations on customer comfort and perceptions was assessed by administering surveys to pilot participants and to a control group of air conditioner load control participants that were not part of the pilot. The control group was called for one system wide event during the summer on September 8 lasting four hours. The pilot participants were called for the same event but also experienced an additional 71 short operations (twice every weekday during the months of August and September). None of the events were preannounced, but customers could opt out during events via web or by calling PG&E.

Because there can be substantial variation across cities and even within cities (e.g., housing vintage, income, etc.), the

Table 2: Estimated reductions by temperature range

	Customers Sent Mail Surveys				Respondents			
Characteristic	Pilot participants (n=)	Control group (n=)	t	p. value	Pilot participants (n=)	Control group (n=)	t	p. value
Energy efficiency rebate in past 5 years	98%	98%	0.16	0.88	99%	98%	0.38	0.71
Low income rate (CARE)	20%	21%	-0.65	0.52	14%	17%	-0.67	0.5
Thermostat device	24%	26%	-0.99	0.32	30%	27%	0.56	0.58
Number of AC units	1.13	1.12	0.73	0.46	1.19	1.14	1.27	0.2
Tons per AC unit	2.75	2.68	0.92	0.36	2.56	2.73	-1.12	0.26
Correlation between monthly consumption and CDD	0.43	0.44	-0.42	0.67	0.40	0.45	-1.03	0.3
Neighborhood average household members	3.43	3.44	-0.69	0.49	3.42	3.36	1.56	0.12
Neighborhood median year home built	1989	1988	2.15	0.03	1988	1988	0.10	0.92
Neighborhood median income	84,869	84,596	0.24	0.81	84,095	85,792	-0.78	0.44
Annual cooling degree days (Base 65 F)	1,623	1,621	0.05	0.96	1,682	1,668	0.23	0.82
Annual PG&E bill	3,827	3,640	1.20	0.23	3,888	3,576	1.01	0.31

control group was selected prior to administering the surveys using two main steps. First, the eligible population of participants in the SmartAC program who did not participate in the pilot was narrowed to those linked to the same weather station and in the same primary city area as customers located in the pilot study feeders. Second, control group candidates were selected based on how well they matched the pilot participants across observable customer and neighborhood characteristics. To do so, we used propensity score matching. This technique requires estimation of the probability customers were part of the pilot feeder population (based on observable characteristics), scores pilot participants and control group candidates, and selects the closest match for each participant (a nearest-neighbor algorithm).

Table 3 compares the control group and pilot participants who were sent the mail surveys to those who responded. In total, of the 814 survey attempts, 454 customers responded, producing a net response rate of 55.8%. Of the 454 respondents, 180 were participants in the AS Pilot and 274 were in the control group. The control and pilot participant groups who were sent the survey were highly comparable. Except for a small difference in the median age of households in the neighborhood, 1988 versus 1989, there were no other statistically significant differences. More importantly, there were no statistically significant differences between the pilot and control group participants that responded to the survey.

Customers were asked about their satisfaction with the utility, with the load control program, how many events they noticed during the summer, thermostat setting, household characteristics, and a number of additional questions.

Among respondents, 91% of the control group and 87% of the pilot participants were satisfied with their relationship with PG&E. On average, the control group and pilot participants rated their experience with the load control program as 7.84 and 7.64, respectively on a scale of 1 to 10, with the majority of both groups reporting that they were overall satisfied with the SmartAC program. About 17% of both the control group and the pilot participants reported that PG&E had turned down their air conditioner during the summer. Survey respondents that reported experiencing an event were asked how many events they experienced. Customers in the control group (who in fact experienced a single event) reported on average 3.23 events throughout the summer, while similar customers in the pilot group (who experienced 69 to 72 events) reported an average of 2.79 events throughout the summer.

None of the small differences in customer satisfaction or perception about the number of events were statistically significant. The findings indicate that repeated short term control of AC units in the PG&E pilot did not affect customer satisfaction, perceptions, or comfort. Put differently, the findings suggest that residential air conditioner loads can be used to provide ancillary services with little or no effect on customer comfort or satisfaction.

VI. Conclusions and Recommendations

The PG&E pilot was a large scale test of the ability to use residential AC loads for ancillary service operations. It produced several key findings:

- Air conditioner load control programs can start resources quickly, typically, within 60 seconds and generally ramp up to full capacity in less than 7 minutes, with roughly 80% of the available demand reduction begin to be delivered in under 3 minutes.
- Air conditioner electric demand patterns can be transmitted in near real time, providing operators information about the resources available and confirmation of demand reductions being delivered.
- The demand reductions observed in the air conditioner end-use data were also observable in the feeder loads; however, this is only true under hot temperatures for feeders with a high penetration of participants in air conditioner load control programs.
- The demand reductions that can be delivered vary by time of day and temperature conditions and communication network signal strength. Systematic test operations can provide valuable information about the variation and help produce better estimates of the magnitude of resources available. It can also help identify areas where the communication network requires reinforcement.
- Repeated short term AC curtailments (15 minutes or less) did not lead to statistically significant differences in customer satisfaction or comfort.

Table 3: Comparison of control group and pilot participants

	Antioch			Fairfield			Fresno I			Fresno II		
Temp (F°)	# of events	Estimated Impact (kW)	% Reduction	# of events	Estimated Impact (kW)	% Reduction	# of events	Estimated Impact (kW)	% Reduction	# of events	Estimated Impact (kW)	% Reduction
75 or less	11	-0.27	-87.1%	11	-0.21	-80.8%	3	-0.05	-83.3%	3	-0.01	-100.0%
75-80 F°	10	-0.26	-63.4%	10	-0.21	-67.7%	2	-0.09	-69.2%	2	-0.03	-60.0%
80-85 F°	12	-0.45	-80.4%	12	-0.39	-81.3%	3	-0.28	-59.6%	3	-0.18	-58.1%
85-90 F°	11	-0.75	-81.5%	11	-0.7	-82.4%	6	-0.28	-52.8%	6	-0.1	-32.3%
90-95 F°	12	-0.93	-80.9%	12	-0.98	-83.8%	21	-0.64	-61.5%	21	-0.54	-58.1%
95-100 F°	10	-1.21	-73.3%	10	-1.31	-80.9%	26	-0.82	-55.0%	26	-0.85	-54.5%
Above 100 F°	2	-1.27	-84.1%	2	-1.44	-82.8%	7	-1.13	-62.1%	7	-1.18	-57.8%
Overall	**68**	**-0.64**	**-79.0%**	**68**	**-0.62**	**-80.5%**	**68**	**-0.66**	**-58.4%**	**68**	**-0.62**	**-55.9%**

Many large AC load control programs exist across the U.S., many of which control tens or hundreds of thousands of AC units. For example, PG&E's program currently controls over 160,000 AC units, and a full scale use of those resources could provide approximately 100 MW of load reduction for ancillary services for most summer days, and upwards of 200 MW for system peaking conditions.

However, several additional steps need to be undertaken to utilize AC loads for grid operations and incorporate them into markets. These include determining rules on how to conduct settlement for ancillary services bid into markets by load control programs. Generally, the bulk of the payments are related to availability, with penalties for failure to deliver the resourced bid in. AC load control is a unique resource for ancillary services in that their capability is variable though highly predictable. In addition, telemetry requirements need to be re-defined so that they provide operators to confirm that AC resources have been dispatched without imposing substantial costs. This likely means relying on samples rather than requiring telemetry of each individual unit. In addition, the processes for delivering specific amounts of resources need to be systematically done so operators can request discrete amount of resources (i.e., partial dispatch of AC resources).

Appendix

Acknowledgment

This project was sponsored by the California Energy Commission Public Interest Energy Research (PIER) Program under contract #500-05-001. The project team gratefully acknowledges the support of Ron Hofmann, PIER Advisor, and Kristy Chew, PIER Contract Manager.

We also thank Michael Alexander, Wendy Brummer and Osman Zezen from PG&E for their support in the implementation of the pilot. In addition, we extend our gratitude to several firms that helped with the implementation of the pilot, including Population Research Systems, Cooper Power Systems, EnergyICT, and GoodCents. Finally, we acknowledge Paul Mangasarian, who provided analytical support for the pilot while it was being implemented.

References

[1] Sullivan, M., J. Bode, P. Mangasarian. (2009). *2009 Pacific Gas and Electric Company SmartAC Ancillary Services Pilot.* Prepared for Pacific Gas and Electric. Available at http://fscgroup.com/index.php/publications/124-2009-pacific-gas-and-electric-company-smartac-ancillary-services-pilot.

[2] Kueck K., B. Kirby, M. Ally, and Rice. (2009). *Using Air Conditioning Load Response for Spinning Reserve.* Oak Ridge National Laboratory. ORNL/TM-2008/227.

[3] Federal Regulatory Energy Commission. (2011). *2010 Assessment of Demand Response and Advanced Metering.* Available at http://www.ferc.gov/industries/electric/indus-act/demand-response/2010/survey.asp.

[4] Eto, J., C. Marnay, and C. Goldman. Lawrence Berkeley National Laboratory. J. Kueck and B. Kirby. Oak Ridge National Laboratory. J. Dagle. Pacific Northwest National Laboratory. F. Alvarado, T. Mount, and S. Oren. PSERC. C. Martinez. Southern California Edison. (2001, January). *An R&D Agenda to Enhance Electricity System Reliability by Increasing Customer Participation in Emerging Competitive Markets.* IEEE Power Engineering Society Winter Meeting, January 28-31, 2001.

[5] Kueck, J., B. Kirby, R. Staunton, J. Eto, C. Marnay, C. Goldman, C.A. Martinez. (2001). *Load As a Reliability Resource in Restructured Electricity Markets.* Prepared for the U.S Department of Energy. http://certs.lbl.gov/pdf/load-reliability.pdf

[6] Kirby, B. and J. Kueck. (2003, November). *Spinning Reserve from Pump Load: A Technical Findings Report to the California Department of Water Resources*. Oak Ridge National Laboratory.

[7] Eto, J., Nelson-Hoffman, C. Torres, S. Hirth, B. Yinger, J. Kueck, B. Kirby. Oak Ridge National Laboratory. C. Bernier, R. Wright, A. Barat, D. Watson. (2007). *Demand Response Spinning Reserve Demonstration.* Ernest Orlando Lawrence Berkeley National Laboratory. LBNL-62761.

[8] Eto, J., C. Bernier, P. Young, D. Sheehan, J. Kueck, B. Kirby, J. Nelson-Hoffman, and E. Parker. (2009). *Demand Response Spinning Reserve Demonstration — Phase 2 Findings from the Summer of 2008.* Ernest Orlando Lawrence Berkeley National Laboratory. LBNL-2490E.

[9] Gifford, W., S. Bodmann, P. Young, J. Eto, J. Laundergan. (2010). *Customer Impact Evaluation for the 2009 Southern California Edison Participating Load Pilot.* Lawrence Berkeley National Laboratory. LBNL-3550E.