Demanding the Most from DR: How Evaluation Can Inform Residential Demand Response Optimization

Alexandra Horne and Noah Fraser, ADM, Reno, NV
Haixiao Huang, NV Energy, Las Vegas, NV

ABSTRACT

Does it matter if your demand response (DR) event starts at 16:10 or 16:20? Which time is likely to yield more kW savings? In this paper, we explore what granular data can say about optimizing DR. For utilities incorporating smart meters into their residential infrastructure, the ability to utilize 15-minute interval data for measurement and verification (M&V) analyses provides the unique opportunity to untangle high-resolution demand reduction (kW) differences across various event start times. If those same utilities run DR events in phases – with start times for various subgroups staggered by 10-minute intervals, for example – the M&V process can easily yield a detailed comparison of the effectiveness of DR events triggered at a variety of slightly differing times.

By drawing on M&V findings from four different smart and programmable-communicating thermostat models over two DR seasons (87 DR events), we assessed the relative efficacy of an assortment of DR event start times. Herein we discuss the value of triggering events at a variety of times, how these trigger times compare to local peak demand, and compare the relative kW reductions from each start time for each device type to provide tangible guidance on when to trigger your DR events. Overall, the results indicated that events triggered one-hour prior to peak demand may achieve the highest demand reduction.

Introduction

A common method for mitigating periods of peak demand is the implementation of a direct load control program, in which a variety of thermostat types are remotely operated via radio, telephone, or internet connections. Through these programs utilities can trigger Demand Response (DR) Events during which thermostat setpoints at participating premises are shifted by several degrees for approximately one to three hours. These programs can be a particularly useful load-shifting strategy during the summer cooling season, as DR events can be an effective means of shifting the air conditioning load away from the peak demand period. However, during the approximately two-hour window of time after a DR event, the increased cooling load can be significant – anywhere from 5 to 20% higher than normal consumption at this time of day. This period, often referred to as "snapback", can significantly impact the grid when enrollment is scaled-up (i.e., many tens of thousands of premises). Moreover, some implementation contractors run pre-cooling strategies that occur prior to DR events, during which thermostat set points may be decreased for one to two hours prior to the event's start to maintain more comfortable conditions within the premise during the event itself. These pre-cooling strategies impact the grid in a similar way as snapback, causing a period of increased consumption as air conditioning systems work to cool spaces to lower-than-normal temperatures.

One strategy for mitigating the effects of snapback and pre-cooling on the grid when DR programs have high participation is to avoid triggering DR events for an entire population simultaneously and to instead run events in phases staggered by 10-minute intervals over a period of one plus hours. While this may not decrease the total consumption of any given premise during either the snapback or pre-cooling windows, this method does tactically distribute the undesirable load impact of snapback and pre-cooling from the entire DR program across a broader period. In this paper, we leverage the implementation of such a phased approach, along with the availability of advanced metering infrastructure (AMI) datasets from smart meters, to evaluate if there is an appreciable difference between the demand reduction achieved by events initiated at a variety of slightly offset times.

Methodology

The analysis methodology described herein was implemented over two program year evaluations to compile the data reviewed in this paper.

Data Acquisition and Sampling

AMI data was provided for demand response event analyses by the utility for about 75% of program participants. Since smart meters record premise energy consumption data per 15-minute interval, we summed these values in to convert to total hourly demand (kW). To qualify for inclusion in the analysis sample, premises needed to maintain enrollment throughout DR season and to not have participated in any other Demand Side Management program offerings that year. Additionally, the analysis sample was limited to premises with only one kind of demand response device installed. For example, if a premise had both a programmable-communicating thermostat (PCT) and a smart thermostat device installed, then it is excluded from the analysis. Conversely, a premise with two PCTs installed was retained for analysis.

Event Data Cleaning

For each event analysis, several additional data cleaning steps were performed to ensure accurate event evaluation. Since all devices provide a way for the customer to override a demand response curtailment event, any premises that may have overridden an event needed to be identified. Overrides are tracked through various data resources, all of which are provided by the utility and their implementers, and any premise that was found as having overridden event curtailment was removed from the event analysis for that day.

Additionally, premises with non-responding devices—devices that do not act in response to the curtailment signal sent, potentially due to system outages or disruptions to internet connectivity— were identified and removed from the event day analysis. Two methods were used to ensure proper identification of non-responding devices and limit Type One errors – a cumulative time series evaluation that identifies a break in slope and a direct screening for a set percent reduction in consumption over the first event hour. Only premises that fail both tests were removed from the analysis.

After cleaning the population of premises based on the above criteria, the event analyses included in this study incorporated a sample of approximately 45,000 single-family residential premises and 5,000 multi-family residential premises on average per event.

Determining the Event-Day Baseline

Following data cleaning, the individual consumption time series for all premises were aggregated by device type, rate classification, and event phase to define distinct “subgroups” for analysis. This approach smoothed out the noise present in individual premise consumption patterns and served as a representative sample of the average household demand.

The utility provided a list of over 62,000 randomly selected single-family premises and over 6,000 randomly chosen multi-family premises for use as a control group for the analysis. Up to five control premises from this population were matched to each DR premise based on propensity score matching utilizing property characteristics such as location (zip code), square footage, and home age. Then for each event, the matched pool of control premises was re-matched to each subgroup using a pseudo-clustering method based on total consumption in the seven preceding non-event weekdays. This method resulted in several matched control group premises per home participating in the DR event.

The consumption time series from the selected control group premises were aggregated to create a baseline consumption time series for each subgroup. Even with these matching steps, there can still be a consistent difference between the event-day participant and control group time series that could lead to a bias in the difference of the two curves. To rectify any disparities, an event-day adjustment ($\alpha_g(t)$) was made the average control group load ($\overline{L_g^c}(t)$):

$$\overline{L_g^c}(t) \rightarrow \overline{L_g^c}(t) \cdot \alpha_g(t)$$

For devices with pre-cooling enabled, the adjustment factor was based on the mean of the ratios of the average treatment group load and the average control group load for the hour before pre-cooling began and the third hour after the event ended:

$$\alpha_g(t) = \ mean(\frac{\overline{L_g}(t_{pre})}{\overline{L_g^C}(t_{pre})}, \frac{\overline{L_g}(t_{post})}{\overline{L_g^C}(t_{post})})$$

For thermostats that do not include pre-cooling periods prior to the event, the adjustment factor was calculated as the ratio of the average treatment group load and the average control group load during the hour prior to the event:

$$\alpha_g(t) = \overline{L_g}(t_{pre}) / \overline{L_g^C}(t_{pre})$$

An example load shape from a selected control group (Baseline) is plotted in Figure 1 alongside the event-day adjusted curve (NormBL) and the load shape from a treatment subgroup with pre-cooling enabled (Treatment) to provide a visual comparison.

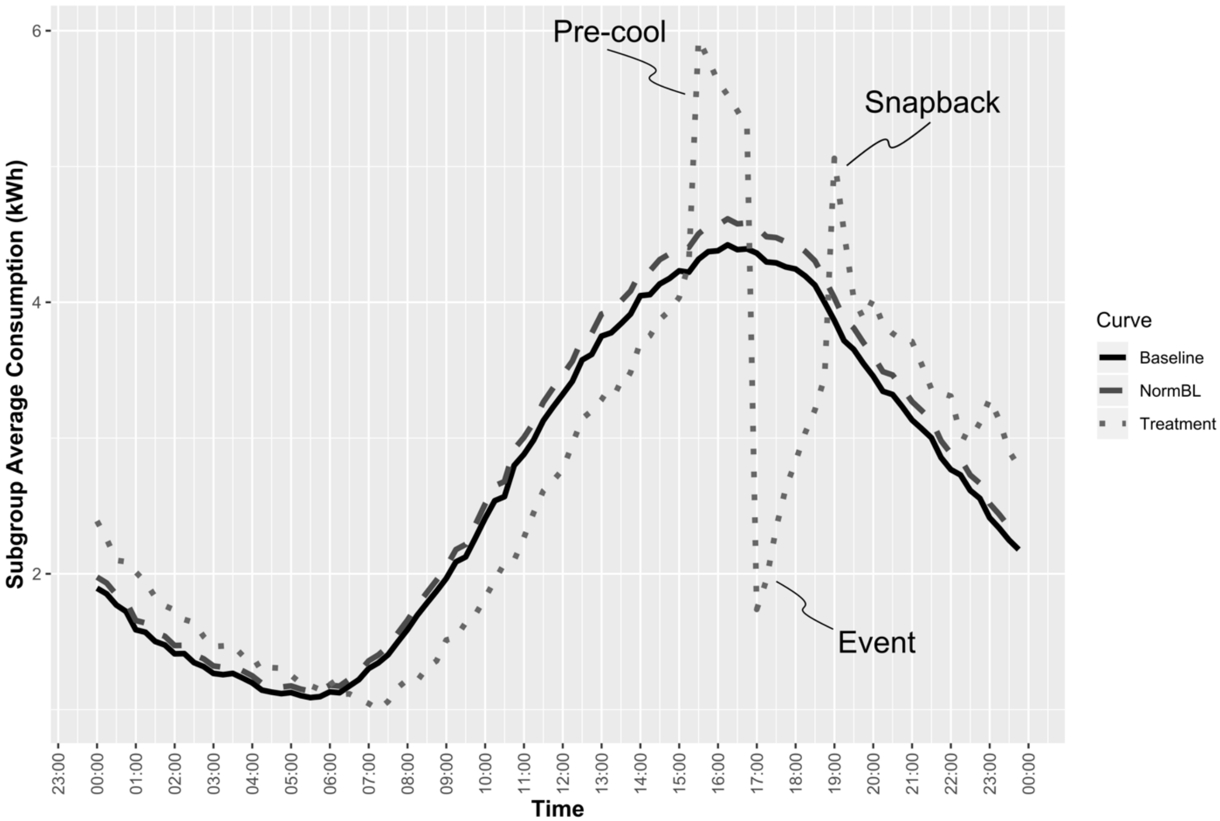

Figure 1. Example Event Day Load Shape

Energy savings related to the demand response event, which began at approximately 16:30 in the example provided above, is illustrated by the heavy consumption drop coinciding with that time and can be considered the area between the event-day adjusted curve and the treatment subgroup curve. Conversely, pre-cool and snapback energy consumption is represented by the sharp increases in consumption plotted by the treatment subgroup curve starting at approximately 15:00 and 19:00, respectively. Smart thermostat optimization algorithms can give the treatment subgroup load shape a shifted appearance but are very common.

Well-illustrated in Figure 1 is also the local peak demand hour, which is commonly one hour before 17:00 during the DR season (June – September). Peak demand hour was determined by ADM during a separate critical peak demand analysis completed for the utility. The DR events discussed in this study ran for two hours following the listed phase start times, aligning broadly with the system peak.

Calculating Demand Reduction

The demand reduction calculated for each subgroup is based on the average load differential for each hour of the DR event. The subgroup kW factor is the upper limit of a subgroup's per-device load reduction capability and is determined by normalizing the maximum load reduction achieved during the event by the ratio of the number of devices per premise for the subgroup. In this paper, only the end-of-line demand reductions are reviewed, and line losses are not included in the computation of the demand reduction achieved.

Weather Normalization

To enable the direct comparison of event results over multiple DR seasons, an additional step of creating weather-normalized kW factors would be typically recommended. Weather data from the nearest local weather station in the American southwest was downloaded from the National Oceanic and Atmospheric Administration and used to calculate the average daily temperature and cooling degree days for each event day. The event-day cooling degree days were then used to normalize the calculated kW factor for each subgroup.

However, as shown in Figure 2 below, there does not appear to be a strong correlation in demand reduction associated with temperature (the apparent increase in kW factor scatter with temperature is an artifact of having more data from hotter days). Therefore, we chose not to normalize the event kW factors to cooling degree days as this data processing step may only serve to blur variability between the demand reduction realized by various event phase start times.

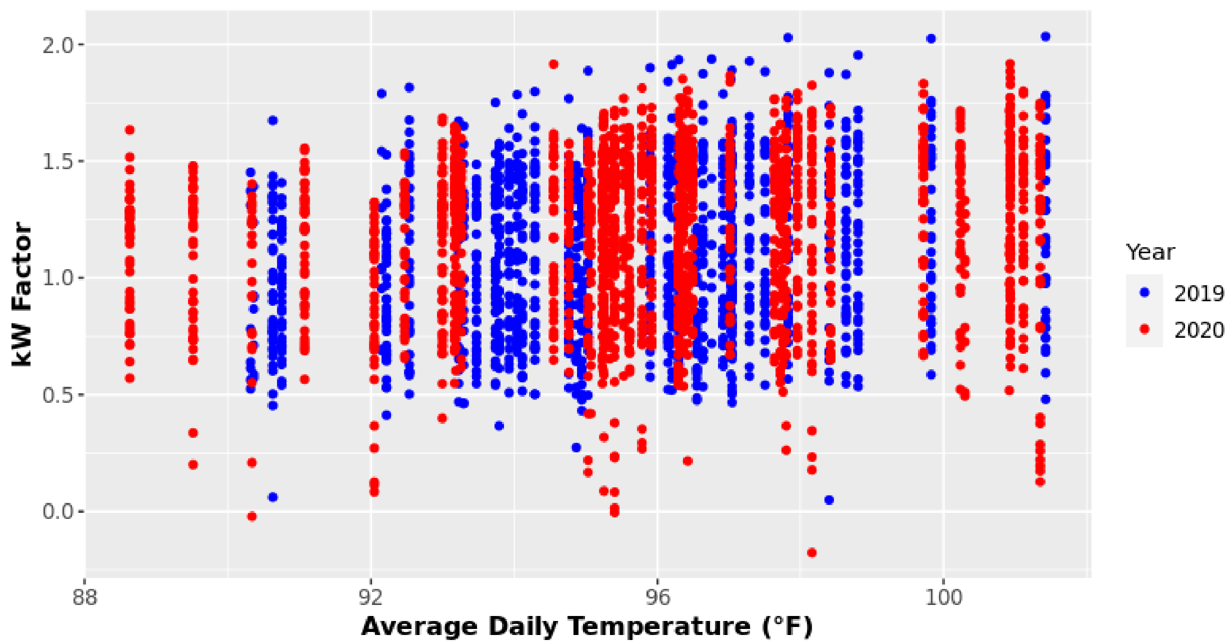

Figure 2. Event Day Temperature vs. kW Factor

Results

Mean kW factors attained across all thermostat devices for event phases starting between 16:00 and 17:10 are presented in Table 1. The number of subgroup events included in the calculated mean values ranged from 302 to 472 analyses. While the mean kW factor determined for various start times varies by only 0.23 kW or less, a Kruskal-Wallis test p-value of << 0.05 indicates statistically significant variability in kW factor across start times. This variability is dampened in the weather-normalized kW factor dataset also presented in Table 1 for comparison.

Table 1. Mean kW Factor by Event Start Time

Start Time[1]	Mean kW Factor	Mean Weather Normalized kW Factor	N
16:00	1.27	0.06	438
16:10	1.14	0.06	420
16:20	1.10	0.05	416
16:30	1.14	0.06	472
16:40	1.18	0.06	472
16:50	1.04	0.05	362
17:00	1.22	0.06	392
17:10	1.07	0.05	302

To further explore any variability present in the data set, kW factors were disaggregated by both thermostat device type as well as by the utility's rate classification. Among the four thermostat models reviewed in this study, there is considerable overlap in event performance across all phase times. As shown in Figure 3, box plots illustrating quartile statistics for each device type show very similar distributions of kW factor results for each event phase start time, though events triggered at 16:00 may be slightly more impactful with higher kW factors on average. Smart Thermostat B overall appears to perform the most consistently across all events, while Smart Thermostats A and C generate slightly more variable results. For the PCT thermostats, the inter-quartile range for event phase start times between 16:00 and 16:50 are slightly higher and more variable than the same metric for events triggered later, however limited data is likely the primary driver for this apparent deviation as there were only a small number of events triggered for this device between 17:00 and 17:10 over the two DR seasons.

Likewise, Figure 4 demonstrates a lack of distinct differences in demand reduction across the spread of event phase start times for both single-family and multi-family premises; however, there is also evidence in this disaggregation that event phases triggered for single-family premises at 16:00, one hour prior to peak demand, may yield slightly higher demand reduction.

To explore the tendency of single-family homes indicating higher savings at 16:00 in more detail, Figure 5 disaggregates the results from single-family premises by device type. For all three smart thermostats installed in single-family premises, event phases triggered at 16:00 do tend to generate slightly higher demand reduction, although for PCT devices this distinction is not as pronounced. Additional Kruskal-Wallis analysis of variance between the mean kW factors across the various event phase start times for each device type installed in a single-family premise also indicate significant deviation in average demand reduction between various event phase start times for every smart thermostat ($p < 0.05$ in all cases) but not for PCT devices (p value of 0.22 when the three events triggered at 17:00 or later are removed from the sample).

[1] Additional event phases were triggered between 15:30 – 15:50 as well as between 17:20 – 18:00, however less than 40 events were triggered at any one of those phase times and therefore they are not included in this review.

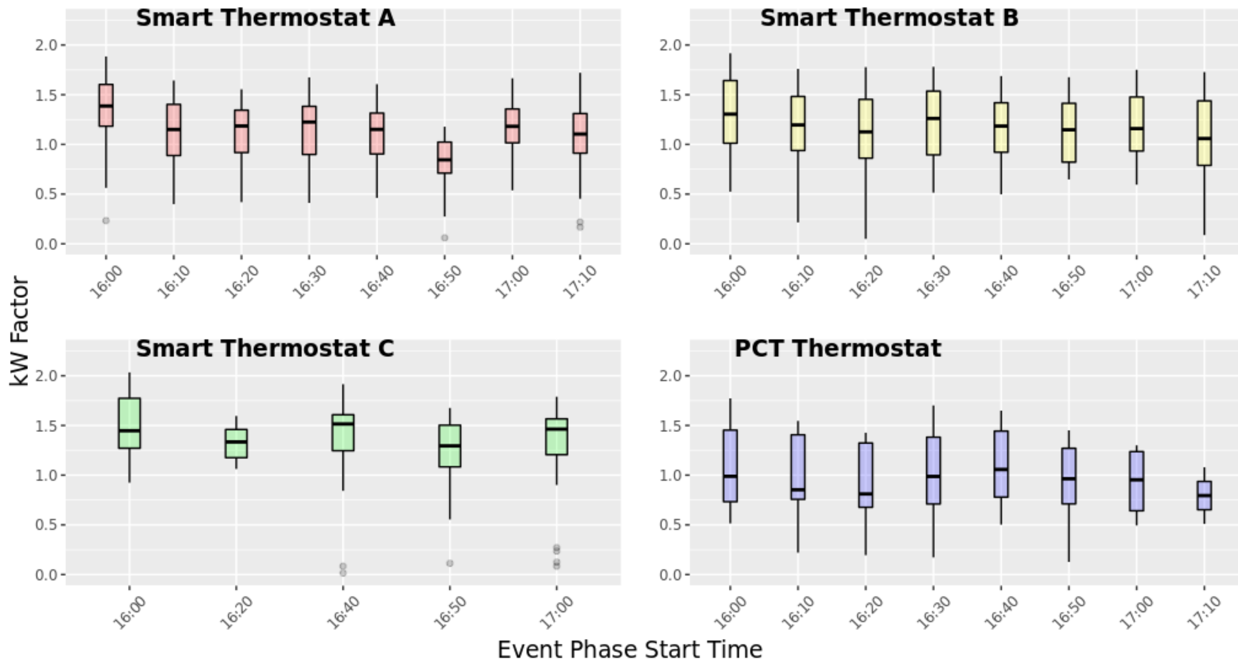

Figure 3. Impact of Event Phase Start Time by Thermostat Device

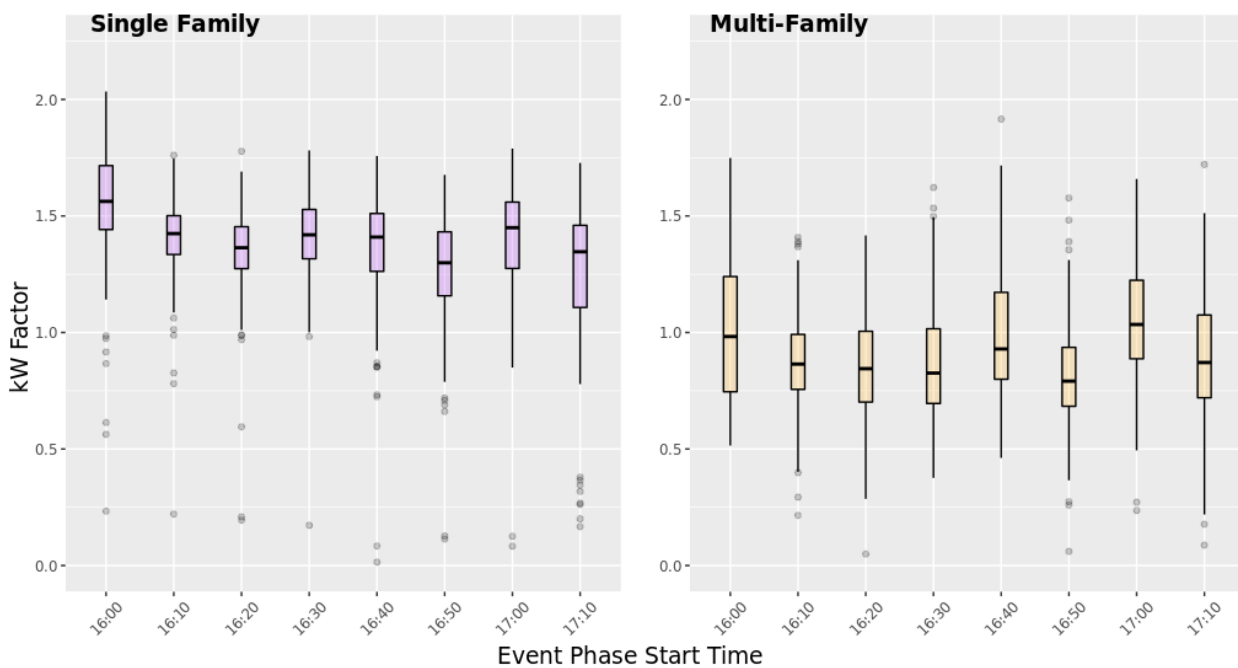

Figure 4. Impact of Event Phase Start Time by Rate Class

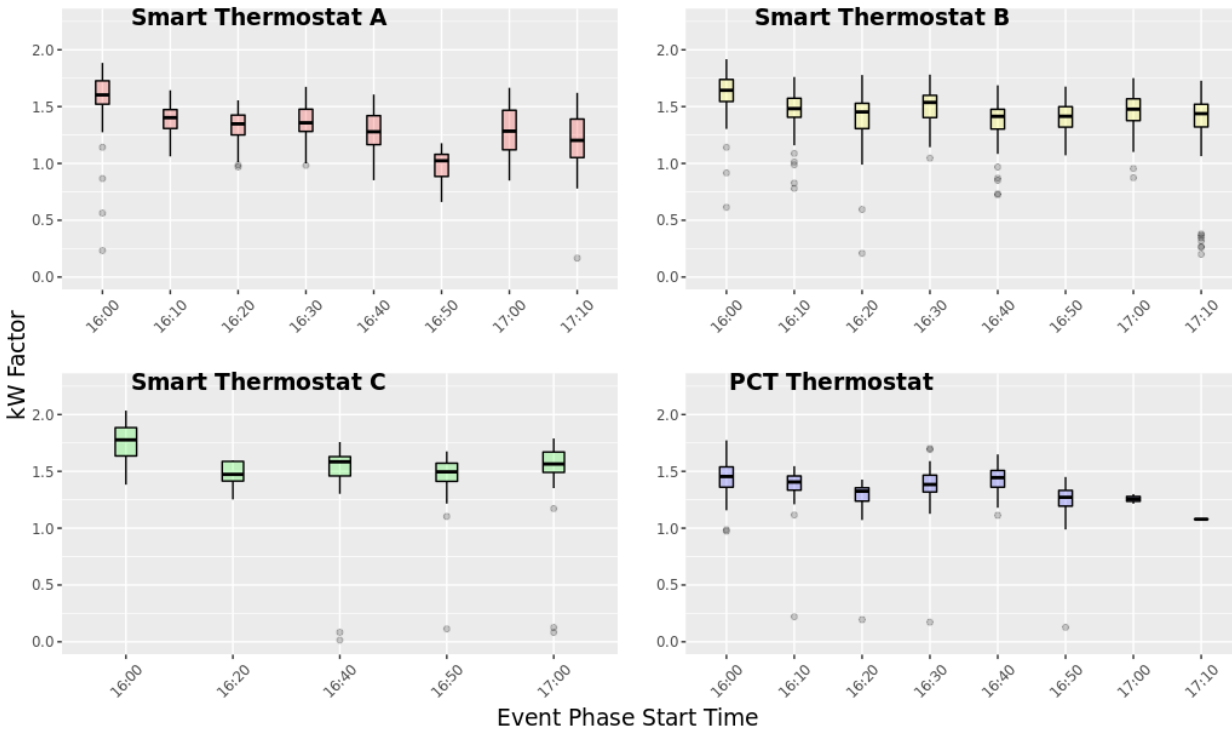

Figure 5. Impact of Event Phase Start Time by Device Type for Single Family Homes

Discussion and Conclusions

The body of data presented in this study collectively revealed that over two demand response program years there were not significant differences in the demand reduction achieved in multi-family premises by load shifting events triggered across eight different 10-minute intervals in the hour preceding the system's peak demand. However, events triggered over the same DR seasons in single-family homes with smart thermostat devices do indicate that demand reduction is slightly higher when the two-hour DR event is triggered one-hour prior to peak demand.

Overall, these findings can be viewed as encouraging for utilities interested in expanding their demand response programs. By distributing participating premises across a range of event start times, the impact of pre-cooling and snapback increased consumption for high-enrollment DR programs can be mitigated without causing a drastic variability to the demand reduction capabilities of the program. Though care must be taken when scheduling phased demand response events to avoid initiating either a pre-cooling or a snapback surge in consumption coincident with the system's peak demand hour, with strategic planning, the utilization of phased demand response event start times can enable enhanced load control capabilities.