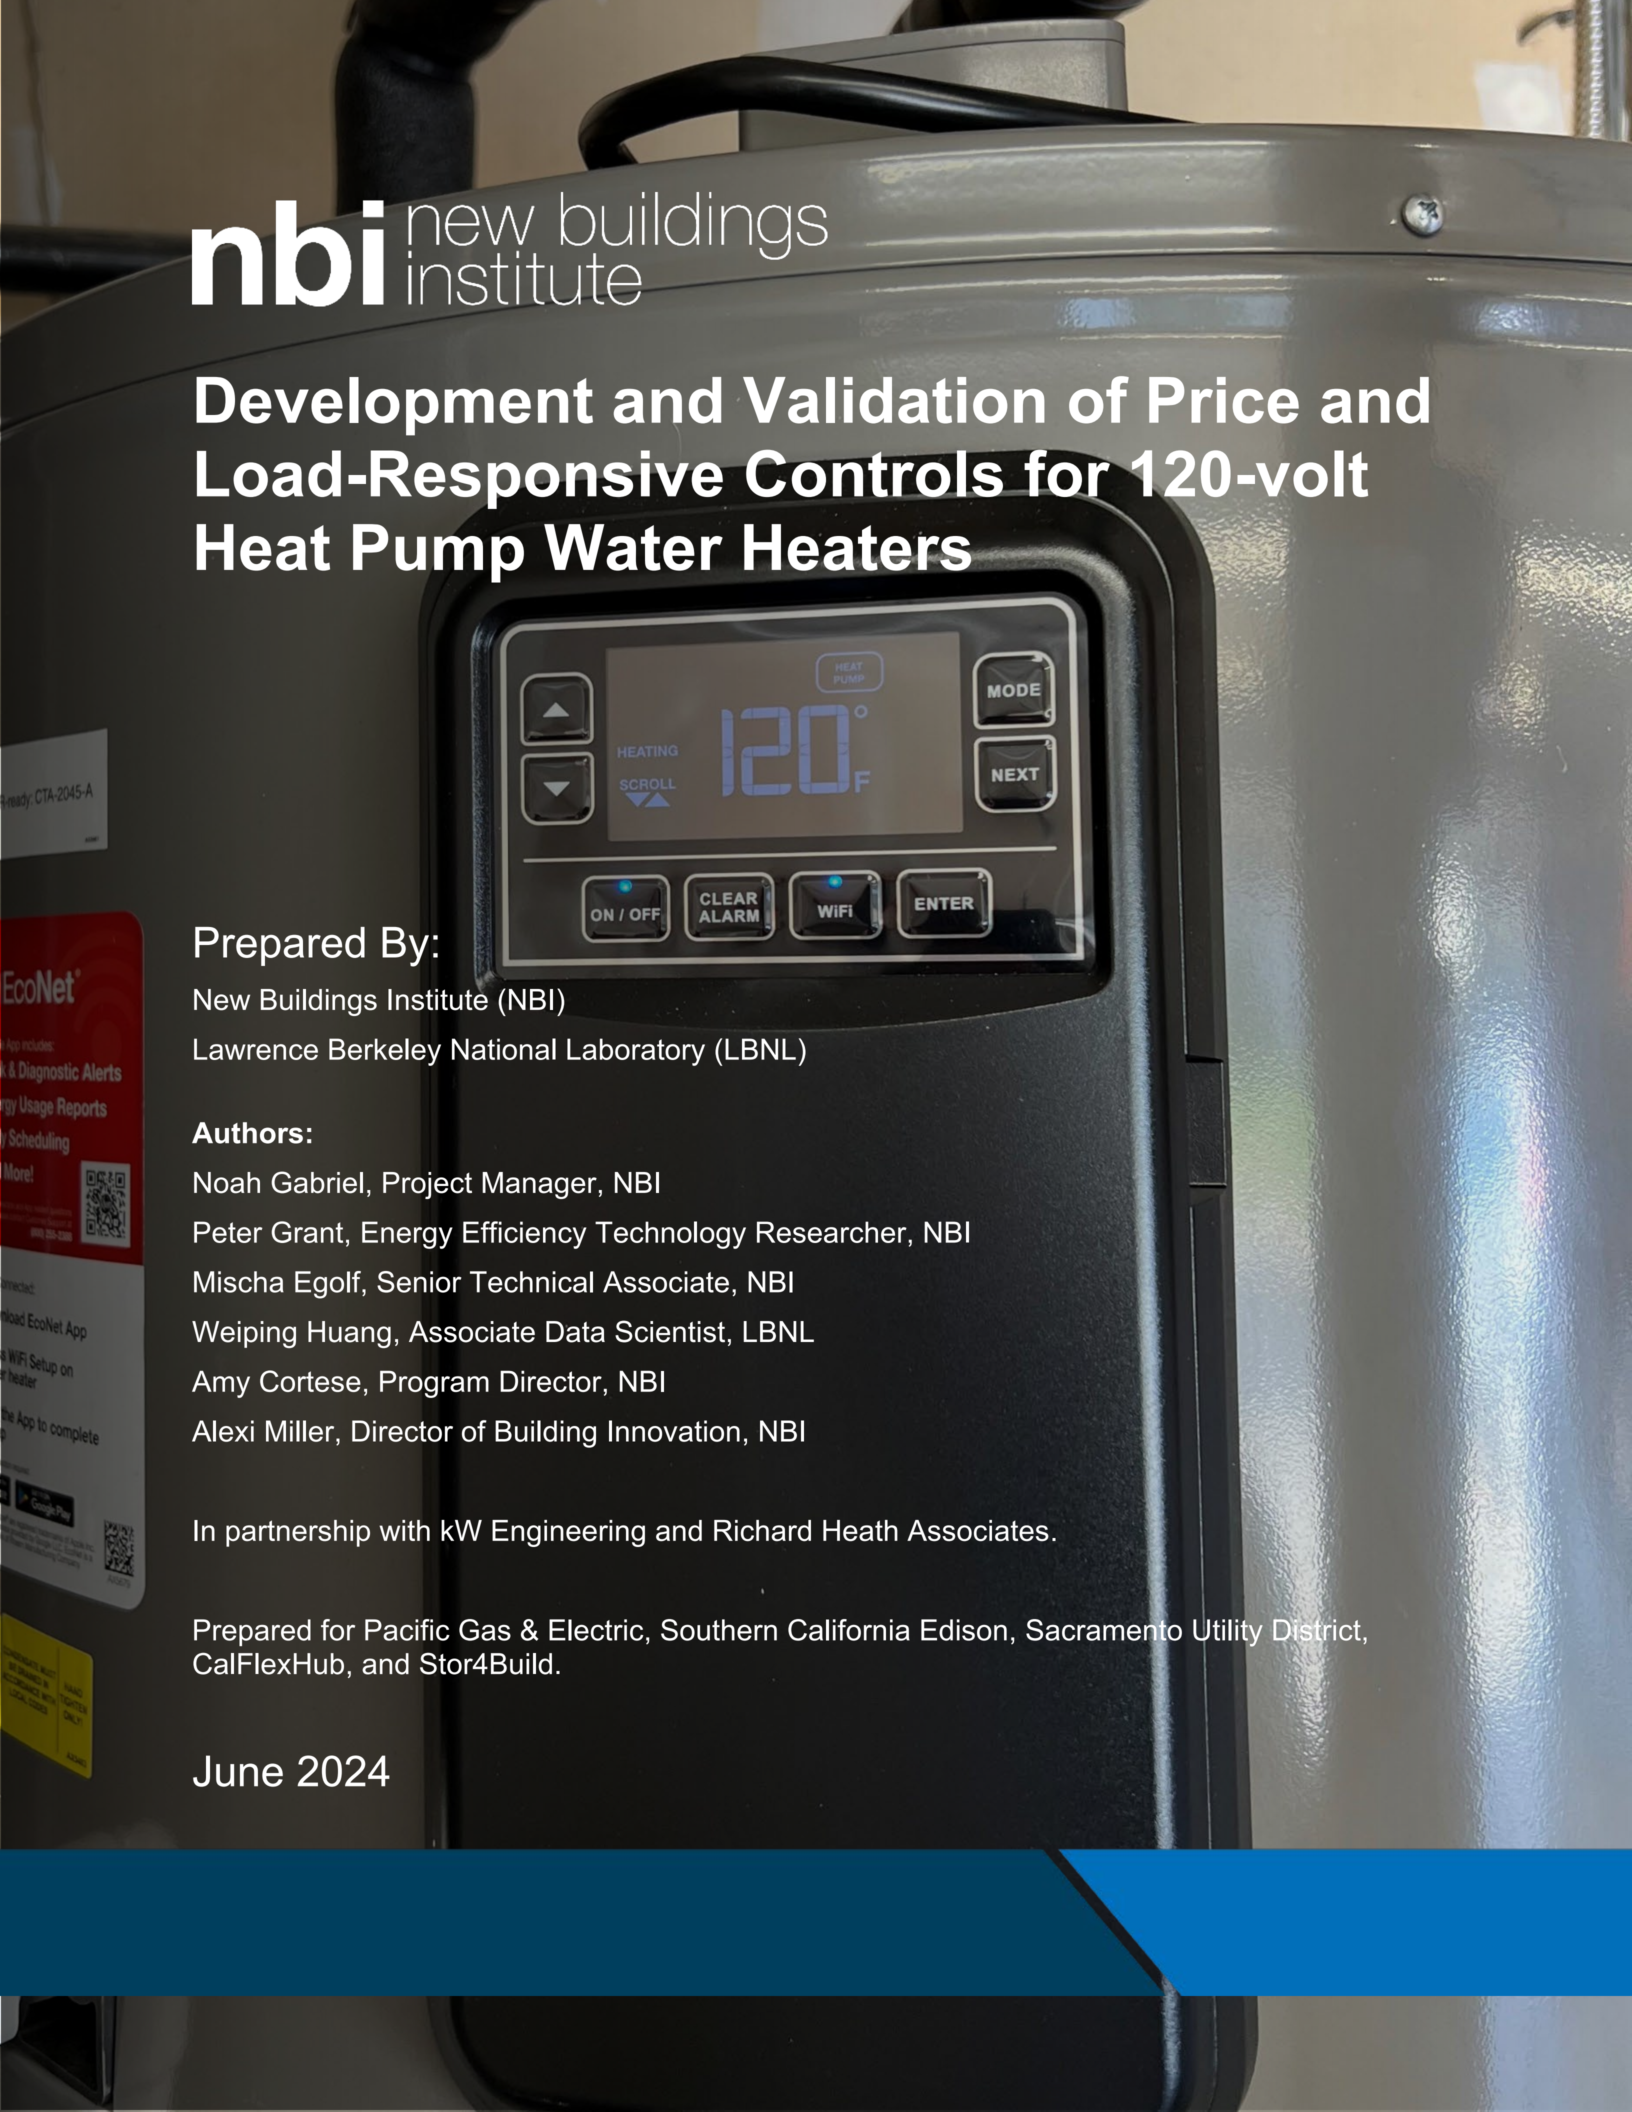

nbi new buildings institute

Development and Validation of Price and Load-Responsive Controls for 120-volt Heat Pump Water Heaters

Prepared By:

New Buildings Institute (NBI)

Lawrence Berkeley National Laboratory (LBNL)

Authors:

Noah Gabriel, Project Manager, NBI

Peter Grant, Energy Efficiency Technology Researcher, NBI

Mischa Egolf, Senior Technical Associate, NBI

Weiping Huang, Associate Data Scientist, LBNL

Amy Cortese, Program Director, NBI

Alexi Miller, Director of Building Innovation, NBI

In partnership with kW Engineering and Richard Heath Associates.

Prepared for Pacific Gas & Electric, Southern California Edison, Sacramento Utility District, CalFlexHub, and Stor4Build.

June 2024

RESEARCH TEAM AND ACKNOWLEDGEMENTS

This study is the result of funding and contributions from a wide range of organizations and individuals.

The following utilities and research groups provided funds in support of the research: Pacific Gas and Electric (PG&E), Sacramento Municipal Utility District (SMUD), Southern California Edison (SCE), Stor4Build, and CalFlexHub.

Research Team:

Noah Gabriel, Project Manager (NBI)
Mischa Egolf, Senior Technical Associate (NBI)
Amy Cortese, Program Director (NBI)
Peter Grant, Energy Efficiency Technology Researcher (LBNL)
Weiping Huang, Associate Data Scientist (LBNL)
Tao Yang, Postdoctoral Researcher (LBNL)
Aditya Kanteti, Student Assistant (LBNL)

Utility Partners:

Kelly Cunningham, Program Manager (PG&E)
Thomas Mertens, Program Engineer (PG&E)
Helen Liu, Senior Program Manager (PG&E)
Charles Kim, Senior Engineer (SCE)
Terry Borges, Project Manager II (SMUD)

Table of Contents

Figures and Tables

EXECUTIVE SUMMARY

New Buildings Institute (NBI) oversaw the installation of CTA-2045-enabled Universal Communication Modules (UCMs) on twelve 120-volt (120V) heat pump water heaters (HPWHs). These installations were spread throughout three utility territories in California.

NBI and Lawrence Berkeley National Laboratory (LBNL) used UCM technology to conduct the first-ever field tests on the load shifting efficacy of these appliances. The study team performed five week-long tests over a five-month period. Four of these tests used the *Load Up* (LU) CTA-2045 signal to increase thermal energy storage ahead of an energy curtailment period, and one used the *Advanced Load Up* (ALU) signal.

NBI and LBNL found that the 120V HPWHs shifted an average of 39.4% of HPWH load away from peak utility demand periods when using the LU signal. This load shift resulted in average utility bill savings of 24.2% of HPWH electricity costs.

The water heaters did not respond correctly to the ALU signal during field testing, presumably due to lack of technological support from the manufacturer. However, LBNL modeling predicts that, with successful ALU conformance and using CalFlexHub electricity prices, these HPWHs could increase the load shifted away from the peak period by an additional 31% and the operating cost reduction by 21%.

BACKGROUND AND INTRODUCTION

Phase I Background & Results

The first retrofit ready, plug-in 120V HPWH models entered the U.S. market in 2021. The low power, 120V design can plug in to existing wall outlets without the expensive panel upgrades and/or home rewiring often required for traditional 240V units. The 120V HPWH is an ideal solution for fuel-switching retrofit applications to replace existing gas-fired tank type water heaters; it is especially well suited for smaller homes with lower hot water demand.

In July 2023, NBI completed the initial phase – referred to as "Phase I" – of the nation's first 120V field study. During the study, the research team collected data from 32 participating sites throughout 13 California climate zones and three utility territories.

This data provided several key insights as reported in the Phase I findings[1]. On average, study participants used 85% less energy to heat water after switching from a gas water heater to a 120V HPWH. Operating costs also fell by more than 50% on average for customers in PG&E and SMUD service territories. Operating costs increased slightly in the winter months for SCE customers where gas prices are low relative to the rest of the California market. The current draw for HPWHs in the study typically peaked below 5 amps, far lower than 240V HPWHs that typically reach 21 amps. This low draw reduces grid strain, which is critical as California moves towards widespread electrification.

Based on customer satisfaction and high efficiency, these innovative water heaters can serve as a plug-and-play decarbonization solution for millions of homes, especially in California where gas water heating is dominant. NBI researchers estimated that 22%-30% of California homes could be directly supported by plug-in water heaters[2]. This estimation was based on the percentage of households that meet the following criteria where 120V HPWHs are best suited:

- 1–4-person occupancy
- An existing propane or gas water heater
- Adequate air space from which to extract heat (450+ cubic feet)

Phase II Background & Scope

Building electrification and the integration of variable renewable energy resources will exacerbate operational and financial challenges for utilities and grid operators, especially during high-demand periods on the grid, unless demand response measures are taken.

---

[1] For information on Phase I of the study, please see the published findings at https://newbuildings.org/resource/plug-in-heat-pump-water-heater-field-study-findings-market-commercialization-recommendations/.

[2] Khanolkar et al. *Plug-In Heat Pump Water Heater Field Study Findings & Market Commercialization Recommendations* (July 2023). https://newbuildings.org/resource/plug-in-heat-pump-water-heater-field-study-findings-market-commercialization-recommendations/

With their ability to use energy flexibly and act as small thermal batteries, HPWHs can help utilities avoid the use of more expensive generating resources and minimize renewable energy curtailment. Through load shifting signals, the HPWH compressor can be told to run (“charge”) when cleaner generation sources like solar are plentiful. With a full tank of hot water, the HPWH can then coast through the peak demand period without using the energy required to run the compressor. Because California investor-owned utilities (IOUs) enroll their customers in time-of-use rates (TOUs) by default, charging HPWHs during off-peak times with cheaper, cleaner electricity offers the potential for customers to decrease their electric bill.

To test the load shifting capabilities of 120V HPWHs in the field, NBI partnered with LBNL and its load flexibility research program, CalFlexHub (CFH). NBI led project management efforts and recruitment while LBNL developed load shifting controls and conducted data analysis. This second stage of research is hereafter referred to as “Phase II.”

Load Shifting Technology

The CTA-2045 specification is an open standard that enables utilities or third parties such as study teams to remotely send demand response signals to a fleet of water heaters and receive data and current water heater status in exchange. To enable load shifting, the study team utilized CTA-2045-enabled universal communications modules (UCMs) manufactured by SkyCentrics. The CTA-2045 standard is in continuous development. The most current version is called CTA-2045B and is the version used by the study team.

CTA-2045B has six basic signals. Three of these signals direct the water heater to curtail energy use: *Shed, Critical Peak Event, and Grid Emergency.* Shed is the least extreme signal and is designed to be used frequently. *Critical Peak* and *Grid Emergency* are reserved for more extreme load reduction events. CTA-2045B also includes a *Baseline* signal, a *Load Up* (LU) signal, and an *Advanced Load Up* (ALU) signal. LU directs the water heater to preheat its tank until the setpoint temperature is reached and is typically used to prepare for a curtailment period. ALU acts like LU, but stores additional thermal energy, usually via a higher setpoint temperature. Note that ALU – which is used to raise water temperature over the typical scalding threshold – is only possible if the water heater contains a thermostatic mixing valve.

CTA-2045 UCMs are rapidly emerging as the de facto demand response hardware for residential water heaters. Heat pump water heaters sold in California are required to come with factory-installed CTA-2045 adaptors. These adaptors allow end-users to simply remove a cover plate and plug in the UCM.

Figure 1. UCM being plugged into a 240V HPWH.
Photo courtesy of Peter Grant.

RESEARCH OBJECTIVES AND QUESTIONS

Research Objectives

The objectives of this project were to achieve the following:

An understanding of hot water consumption predictability: This project evaluated how well machine learning (ML) algorithms can learn and predict participants' hot water draw. Because water heating use varies based on occupant behavior, an algorithm that can cluster similar households and estimate hot water draw is essential for shifting energy use without triggering runout events.

A methodology for identifying hot water use patterns without a flow meter: Water heater installations outside of this study are unlikely to include flow meters. Thus, it's critical to develop advanced control strategies based solely on participant electricity consumption (as a proxy for hot water use). The UCMs used in this study provide data on participant electricity use. LBNL attempted to leverage this electricity data to create and deploy controls as part of the technology transfer component of this study.

A custom load shifting strategy for 120V HPWHs: Previous field studies covering HPWH load shifting focused on 240-volt (240V) or commercial-scale HPWHs. Because 120V HPWHs typically have no electric resistance heating and utilize mixing valves, successful load shifting requires a tailored algorithm. This project combined simulation modeling and real participant data to develop these necessary, customized algorithms.

Research Questions and Approach

Typical HPWH load shifting applies a single signal schedule, based on engineering judgement, to all HPWHs in a fleet. This approach can reduce peak period electricity consumption but may not minimize HPWH operating costs for two key reasons[3]. First, every home has a unique hot water and electricity consumption profile; thus, each requires different amounts of load shifting at different times of day. Second, the economic incentives for load shifting vary with different electricity rate structures and seasons; therefore, a one-size-fits-all approach is unlikely to succeed.

Prior work shows the potential to improve HPWH load shifting performance by customizing the signal schedules to different occupant behavior and electricity tariff structures. In a 150-household study conducted in the Pacific Northwest, PNNL increased the load Shed during the evening peak by 16% by customizing signal schedules for groups that showed low/high electricity use during the morning and evening[4]. In a separate simulation, LBNL showed the potential to reduce the operating costs of HPWHs by 29% and peak period consumption by 80% when grouping HPWHs by daily electricity consumption and customizing signal schedules to each group[5].

LBNL's prior methodology customized signal schedules based on cumulative daily electricity consumption. However, grouping by daily electricity consumption does not provide insight into what *time* the electricity is consumed, which is critical when optimizing for time-of-use (TOU) rates. To refine this approach, LBNL analyzed the hot water draw and electricity consumption data from Phase I to create typical behavior profiles for each home. LBNL then created new groups for the fleet[6] of HPWHs by identifying HPWHs with similar typical load profiles and generated new signal schedules with this updated grouping method. Based on this literature and prior work, the study team sought to answer the following questions.

---

[3] PG&E. *Evaluation of Unitary Heat Pump Water Heaters with Load-Shifting Controls in a Shared Multi-Family Configuration* (May 2022). https://www.etcc-ca.com/reports/evaluation-unitary-heat-pump-water-heaters-load-shifting-controls-shared-multi-family

[4] Manesseh et al. *Nontargeted vs. Targeted vs. Smart Load Shifting Using Heat Pump Water Heaters* (2021). https://www.mdpi.com/1996-1073/14/22/7574

[5] P. Grant, M. Stuebs, B. Nordman. *Stabilizing the Grid and Reducing Utility Bills Through Price-Responsive Controls for Heat Pump Water Heaters.* 2023 ASHRAE Summer Conference (2023). https://eta-publications.lbl.gov/publications/stabilizing-grid-and-reducing-utility

[6] Term used to denote all HPWHs under the study team's control.

Table 1: Research questions and approach

Research Question	Approach
1. How much can price- and load-responsive controls reduce occupant electricity bills compared to: a) No load shifting controls b) Load shifting protocols currently used in demand response programs	• Collected baseline (non-testing week) data. • Evaluated operating cost of each HPWH when sending 1) A “LU” signal 1 hour before the high price period, and 2) When sending a Shed signal during the high price period. • Evaluated operating cost of each HPWH when sending the signal schedules customized to different operating patterns.
2. How can study findings be leveraged to optimize load shifting at installations without detailed monitoring?	• Attempted to create HPWH groups based on electric load patterns using data from SkyCentrics universal control modules.
3. How can price- and load-responsive controls be customized and transitioned to deployment programs, such as PG&E’s WatterSaver?	• Continued leading the Advanced HPWH Load Shifting working group, providing communication between LBNL, PG&E, and SCE. • Plan to share development of LBNL’s user-friendly Python scripts (toolchain) for identifying signal schedules (to be developed under DOE Stor4Build funding)
4. How repeatable is occupant hot water consumption, and the resulting HPWH electricity demand curve?	• Evaluated the root mean squared error (RMSE) of hourly hot water and electricity consumption for each HPWH in the study. • Tested the ability of several ML forecasters to predict hot water and electricity consumption for each HPWH using variables such as number of occupants, previous hot water use, time of year etc.
5. How do the controls need to be modified to accommodate the day-to-day variation in occupant behavior?	• Updated grouping methods to better capture occupant behavior in the homes. Prior grouping methods used only the total kWh consumed each day, and the new methods also consider the time of electricity consumption. More details are available in Appendix A. • Evaluated performance of new grouping methods.

METHODOLOGY AND APPROACH

In 2023, NBI recruited participants from existing sites, arranged for delivery and installation of the SkyCentrics hardware, and established a data management process for Phase II. Between January and June 2024, NBI and LBNL conducted five week-long load shifting tests, analyzed the results, and synthesized study findings. The sections below detail the work completed in these areas.

Recruitment Process

Phase II participants came from the pool of 32 existing Phase I participants but were limited by the number of eligible manufacturers. While three manufacturers participated in Phase I, only one manufacturer offered the capabilities necessary to execute Phase II. Of the 32 original participants, 26 were eligible for Phase II due to this manufacturer constraint. Of these 26 customers, 18 were in IOU territory and represented the initial recruitment pool. The remainder were installed in SMUD territory and were recruited following the first two load shifting test weeks.

In October 2023, NBI contacted the 18 eligible IOU customers with an invitation to participate. The invitation email template was approved by LBNL's internal review board (IRB). Three customers requested a call to learn more before signing the electronic consent form. Ultimately, ten IOU customers opted into the program by signing the secure electronic consent form. NBI researchers mailed these ten participants a SkyCentrics UCM (the only new hardware required for Phase II participation) and a set of instructions for do-it yourself (DIY) installation developed by LBNL and NBI.

In March 2024, NBI contacted the remaining eight eligible participants, all of whom were in SMUD territory. Of those eight, two opted into Phase II.

The study team gave participants who opted out of Phase II the option to leave the Phase I monitoring equipment onsite without load shifting or installing a UCM. This way, said participants continued to provide valuable data about hot water use habits. All participants who were given the secondary option consented. This monitoring data ensured that at least one full year of data was collected at a variety of sites.

With consent from the Phase I funders, NBI shared this anonymized data with national laboratories as part of a public, nationwide repository[7] that can be used by researchers to better understand hot water heater deployment and energy use. It has also provided valuable insights into hot water runout events as part of form factor research being conducted by Oak Ridge National Laboratory (ORNL).

Data Management

Once participants successfully installed their UCM, their HPWHs were connected to SkyCentrics' data cloud. NBI and LBNL used this platform to schedule load shifting signals and monitor individual HPWH status (online vs offline). The team planned to use this platform for

---

[7] The Heat Pump & Heat Pump Water Heater Field Database can be accessed here: https://heatpumpdata.energy.gov/

detailed data exports, which could be cross-referenced for accuracy with the Phase I monitoring data from the corresponding time period. However, as described in the *Data Access and Quality Control* section below, the team encountered challenges with the SkyCentrics platform that limited this analysis.

Repeatability of Hot Water Consumption

The research team analyzed the hot water and energy consumption dataset from Phase I to evaluate the repeatability of occupant hot water consumption behaviors. This knowledge was intended to provide insights into how behaviors change from day to day and how to transition the findings from this project to larger utility programs such as PG&E's WatterSaver or SCE's Smart Shift.

LBNL evaluated the repeatability of domestic hot water (DHW) consumption in two ways:

- **Statistical analysis**: LBNL generated box and whisker plots showing the DHW during each hour of the day. Highlighting the range of different consumptions at the same time of different days demonstrated the different loads that the HPWH must satisfy and provided insights into how those loads would cause varying electricity demand curves.
- **Machine learning forecasters**: LBNL applied several machine learning forecasters to the Phase I data set, testing their ability to forecast DHW consumption based on previous data. This analysis evaluated the ability to base controls on forecasted DHW demand in scenarios where flow measurements are available.

Load Shifting Algorithm Iteration

The information below provides a high-level summary of the load-shifting algorithm development methodology; more details can be found in Appendix A.

Creating groups of HPWHs with similar behavior is one of the keys to improving load shifting performance, as it enables optimizing load shifting controls to behaviors. LBNL's approach is to 1) perform simulations on a fleet of 148 HPWHs using monitored 24 hour draw profiles, 2) create groups of HPWHs that have similar simulated electricity consumption patterns, 3) identify the optimal control solution for each group of HPWHs, then 4) deploy the controls to HPWHs in the field by identifying the group to which the field HPWH's baseline data most closely matches.

The 148 HPWHs in the fleet leverage the daily hot water draw profiles from Title 24, providing a broad array of realistic HPWH baseline electric load curves from real occupant hot water consumption data. Creating these baseline load curves enables creating groups of HPWHs with similar baseline electricity consumption and customizing load shifting controls to these different behaviors.

Under a parallel DOE Stor4Build project, LBNL improved upon the existing algorithm developed through past research by 1) improving the utilized grouping method to be more sensitive to time of day, and 2) developing a method for adding HPWHs in the field to groups. The original LBNL method grouped HPWHs by average energy (in kWh) consumed each day but did not effectively separate HPWHs by the *time of day* they consumed electricity. LBNL created a new grouping method sensitive to both total and temporal energy use.

Deploying this improved load shifting algorithm in the field then required a method for identifying the optimal load shifting schedule for HPWHs that are installed in the field. LBNL's algorithm does this by 1) identifying the baseline electricity consumption for the site, 2) and comparing the baseline electricity consumption to the HPWHs in the various groups to identify the group with baselines that most closely match the baseline of the field HPWH. The field HPWH is then sent the load shifting schedule for the simulated group that most closely aligns with the actual load curve for that HPWH.

Load Shifting Scheduling

Once the water heaters were placed into groups, LBNL developed a load shifting control schedule for each group. This process required the following three steps:

1. **Reconfiguring the model**: LBNL's prior schedule development work focused on 240V HPWHs, which include backup electric resistance elements. Since this project focused on 120V HPWHs, LBNL updated the model to emulate HPWHs that do not have resistance elements.
2. **Performing the simulations**: LBNL simulated thousands of signal schedules to choose the signal schedule with the lowest operating cost for a given rate structure.
3. **Evaluating comfort impacts**: The backup electric resistance elements in 240V HPWHs activate when there's a high chance of a cold-water event. Since operating those elements is expensive, LBNL's cost minimizing controls for 240V HPWHs avoid approaches that activate the resistance elements, thereby avoiding runouts. Since the 120V HPWHs in this study did not have resistance elements, the toolchain does not have the same feedback and risks causing runouts. LBNL analyzed the output data to ensure that the signal schedules did not yield cold water events.

LBNL developed signal schedules for multiple test cases – described below – which all sought to minimize operating costs under varying price schedules. Price schedules tell consumers how much they will pay per unit of electricity at different times of the day, days of the week, or season of the year. Generally, price schedules encourage consumers and businesses to use electricity when marginal costs for the electricity provider are lowest.

The test cases represent signal schedules for two CFH Highly Dynamic Price (HDP) schedules[8] and three current utility TOU price schedules, as well as two different assumptions about the capabilities of the 120V HPWHs. Utility TOU price schedules were chosen for their dynamic nature, or to assess HPWH load shifting capability for a certain time of day. They are not reflective of average utility prices across all customers. These test cases are:

Table 2. Test cases used in the study

Test Case	Rationale
CFH WinterHDP, No ALU	Due to uncertainty about 120V HPWHs' capabilities, this test case assumed they cannot respond to the CTA-2045-B ALU signal.
CFH WinterHDP	Assumes that the 120V HPWHs can respond to CTA-2045-B *ALU*. This test identified the 120V HPWH's response to *ALU*, and demonstrated the improvement of performance when ALU is available.

---

[8] See https://calflexhub.lbl.gov/calflexhub-affiliate-program/ for further information.

CFH SummerHDP, No ALU	This schedule showcased the controls' ability to respond to different grid conditions (in this case, in the summer vs in the winter).
SMUD, Summer, No ALU	Current SMUD TOU rates, which incentivize shifting load to a midnight – noon window.
SCE TOU-D-5-8PM, Winter, No ALU	Current SCE TOU rates. schedule features an off-peak period from 8 AM to 5 PM, a peak period from 5 PM to 8 PM, and an off-peak period from 8 PM to 8 AM. Because the lowest price is from 8 AM to 5 PM this price schedule provides an incentive to shift load to the solar peak.
PG&E EV2-A, Summer, No ALU	This price schedule represents the case of a home with an EV charger, which is rapidly becoming more frequent. To incentivize shifting load to off peak periods, EV rates feature very high prices during the peak period.

A visual representation of the hourly price schedules is shown below.

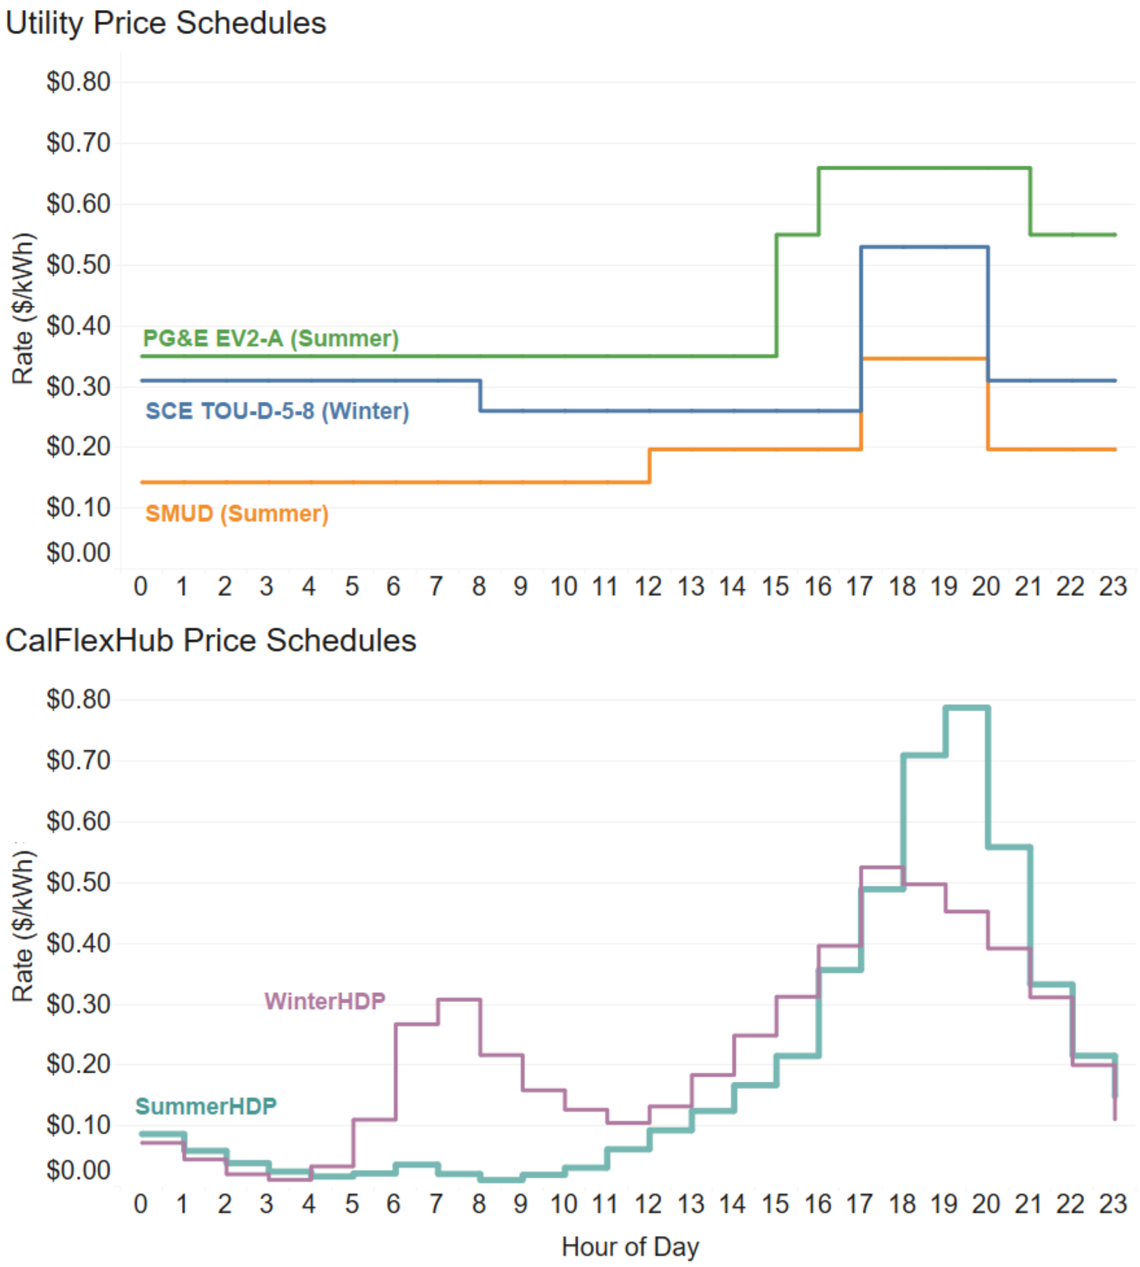

Utility Price Schedules

CalFlexHub Price Schedules

Figure 2. Electricity prices used in the study

After LBNL developed the load shifting schedules for each CFH and TOU price schedule, the study team input the LU and Shed requests into the SkyCentrics platform. Because the HPWHs were grouped by hot water draw profiles, the study team had to manually enter each Load Up or Shed request and specify which group it applied to, as well as the duration of the event and the number of days it would repeat. This process took NBI and LBNL researchers about two hours per new testing schedule for the ten water heaters in the study. Figure 3 displays the SkyCentrics interface for creating a new load shifting events i.e., signals.

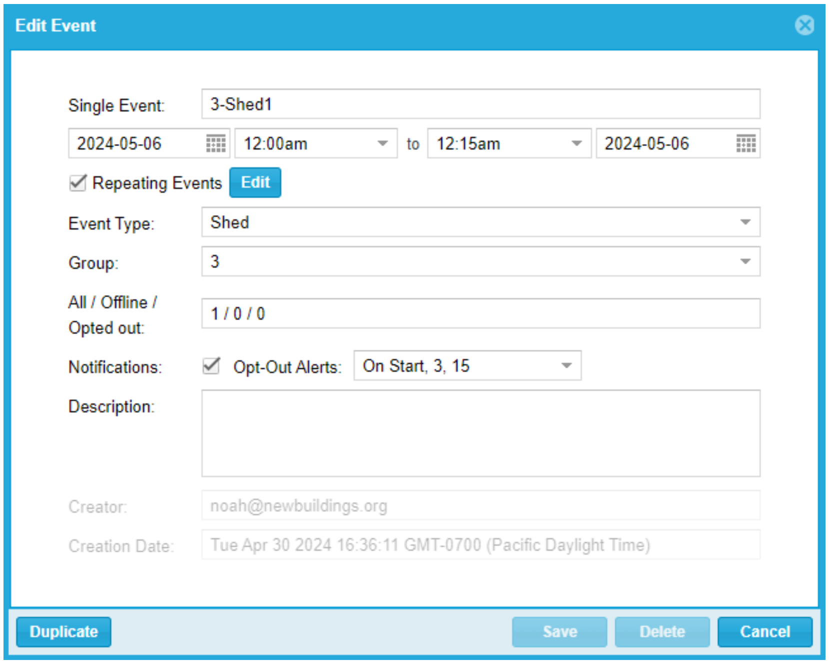

Single Event: 3-Shed1
2024-05-06 12:00am to 12:15am 2024-05-06
☑ Repeating Events
Event Type: Shed
Group: 3
All / Offline / Opted out: 1 / 0 / 0
Notifications: ☑ Opt-Out Alerts: On Start, 3, 15
Description:
Creator: noah@newbuildings.org
Creation Date: Tue Apr 30 2024 16:36:11 GMT-0700 (Pacific Daylight Time)

Figure 3. SkyCentrics Edit Event menu

RESEARCH FINDINGS

Recruitment Challenges & Lessons Learned

Participant Opt-Out Rationale

Only 12 out of 26 potential customers (less than 50%) opted in to Phase II of the study. Reasons for declining the option to participate in Phase II were diverse. For example, one Phase I participant had home solar and believed that no further decarbonization of their HPWH's electricity source was necessary. A common concern surrounded hot water availability, and whether load shifting signals would result in hot water runouts. Some potential participants did not see the value in the study, or believed remote control of their water heater was invasive. This highlighted the importance of communicating how the customer will benefit from load-shifting, the low likelihood of runouts, and how program participants can adjust their habits or level of participation if hot water runouts do occur. It's important to note that the study team only used the Shed signal for curtailment events. Using *Critical Peak* or *Grid Emergency* would be more likely to trigger a cold-water event.

Installation of the UCM

The DIY installation approach greatly reduced study costs by eliminating the need for onsite plumber or electrician visits. However, it also posed some challenges:

1. DIY installation of the SkyCentrics UCMs, while relatively fast, was still an additional ask of participants.

   a. The level of difficulty was a common concern during recruitment. Reassuring participants that help would be provided along with clear instructions was helpful for alleviating concerns.
   b. We recommend that future utility programs record a video of the simple installation for customer use. Doing so would provide clarity for the participants on the exact level of effort and reduce the number of calls/emails needed to answer questions.
2. SkyCentrics' existing installation instructions were made only for 240V HPWHs, so the study team had to adjust them for 120V HPWHs. This process involved speaking with SkyCentrics and the HPWH manufacturer directly and completing a test install to capture photos.
   a. Due to a lack of clarity about whether the 120V HPWHs came pre-wired for CTA-2045 modules[9], we could not precisely tell participants how difficult installation would be. Luckily for this study, all HPWHs were pre-wired. However, verifying the factory-installed capabilities of all participating water heater makes and models will be important for advising customers needing to install UCMs as part of larger utility-run programs.
   b. We recommend that study teams in the future complete bespoke installations before opening communication with participants, especially because manufacturer instructions may lack details or installation edge cases. For instance, one participant's water heater was so close to the ceiling that we could not be sure the UCM would fit.

Failed Installations

Two participants were unsuccessful in their installation attempts, despite assistance from NBI. The NBI study lead confirmed that these two participants correctly followed the installation instructions, and even walked one participant through an installation over video call. However, the two UCMs were offline according to the SkyCentrics website, despite showing a blue light (i.e., Wi-Fi connection successful).

The failure of 17% of the installations suggests that a deeper dive into possible reasons for module failure is necessary ahead of UCM deployment in utility-scale demand response programs. This investigation is especially necessary for DIY programs where user installation errors are expected.

Customer Attrition

Lastly, one participant requested to opt out of the study after receiving the UCM but did not respond to NBI's inquiry into the reason. NBI provided the participant a prepaid UPS shipping label to return the UCM. While this participant was cooperative after opting out, the potential for equipment loss should still be considered when evaluating load shifting program costs.

---

[9] We learned that some 240V HPWHs manufactured before a certain date did not come pre-wired from the factory, because manufacturers expected that the water heater installer would complete this step.

Repeatability of Hot Water Use

Understanding the repeatability – and therefore, predictability – of hot water consumption was a key research objective of this study. The more predictable, i.e., less variable, a home's hot water consumption is, the easier it is to implement a load shifting schedule.

Figure 4 presents the hourly hot water consumption across weekdays from a sample household from Phase I.

While hourly variation over the day is quite high, the data implies that the variation across time periods may be more repeatable. For instance, occupants in Figure 4 may shower between 6-9 AM every day, but slight shifts in the exact timing within that period creates variation in the data.

However, since the HPWH storage tank acts as a thermal battery, the energy load curve may not reflect the exact time hot water was drawn. Further analysis should explore the ability of the storage tank to smooth electric load curves and account for day-to-day variation.

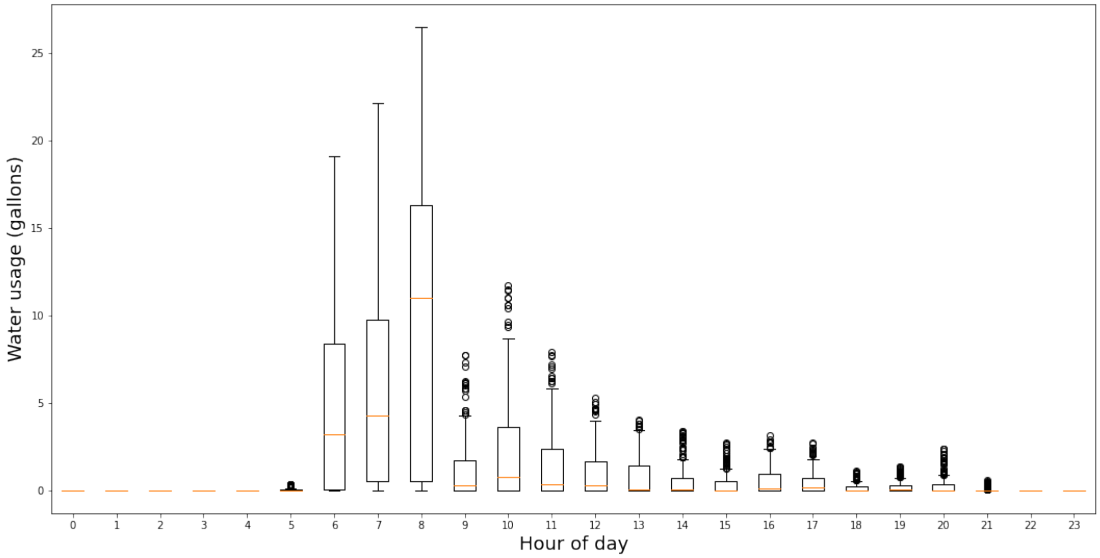

Figure 4. Sample plot showing statistical evaluation of variation in DHW consumption for a single site

Using a machine learning forecaster (a type of predictive algorithm), LBNL forecasted 24-hour-ahead hot water consumption for each site using baseline data. As mentioned above, hourly hot water draws are quite variable, which made accurate prediction difficult.

The machine learning forecaster yielded mixed results due to the variations in DHW consumption. Figure 5 shows a 24-hour forecast of cumulative DHW consumption for a single site on a day when the machine learning forecaster performed comparatively well. The blue bars show the forecasted consumption while the green bars show the actual measured cumulative consumption for that day. Over the 24-hour forecast, the machine learning algorithm predicted 77.35 gallons of hot water consumption compared to 77.43 measured gallons, a different of 0.11%

While the machine learning forecaster accurately predicted total cumulative daily consumption at the end of the 24-hour forecast, the forecasted consumption timing is less accurate. The forecaster overpredicted consumption during the mid-day period before underpredicting during the evening peak. The inaccuracy in timing of the forecast limits the ability of the tool for predictive controls of HPWHs, as this discrepancy could lead to the control either overcharging the HPWH at some times of day or undercharging, leading to a cold-water event at other times of day. In the case of 120V HPWHs, the underprediction during the evening peak would probably cause a cold water event. In the case of 240V HPWHs, the consumption would probably cause electric resistance element use. An element leveraging the consumption forecast would not be able to predict the resistance usage and would not be able to shift load to avoid it.

This challenge appears to be quite frequent when using machine learning algorithms to predict domestic hot water consumption. Domestic hot water consumption does not appear to be correlated with variables such as day of the week or outdoor temperature, instead occurring fairly randomly throughout the day, making it hard for the algorithms to predict. The result is that the algorithms often approach the average consumption during each hour of the day, which minimizes the mathematical error but misses the timing. LBNL tested 12 different machine learning algorithms and found this to be a consistent result.

Furthermore, generating these forecasts requires a detailed understanding of hot water consumption at each site which requires flow meters. Since water heaters do not include flow meters, and adding them would represent an additional expense, the data set required for this approach is probably not viable at scale.

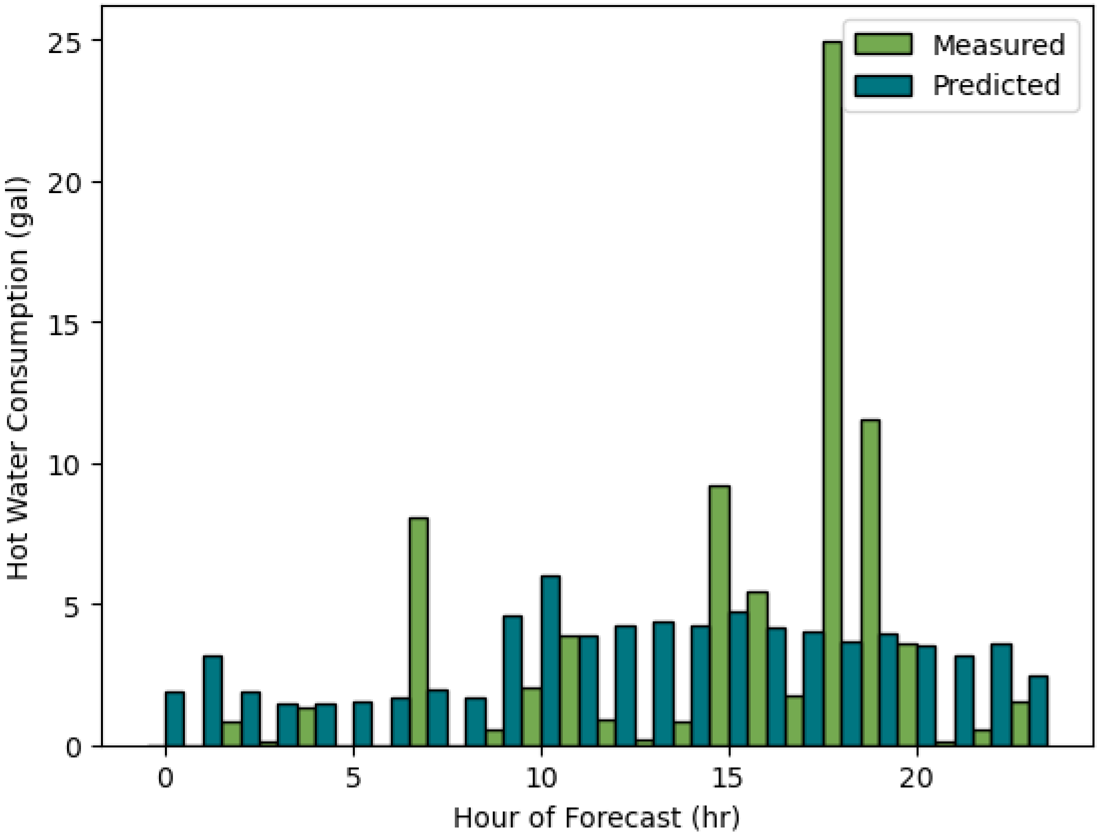

Figure 5. Example 24-hour forecast of predicted vs actual cumulative DHW consumption for a single site

Data Access and Quality Control

One key research objective of this study was to optimize load shifting controls without the use of detailed monitoring equipment. As a result, collecting baseline electricity data through the CTA-2045 module was vital. SkyCentrics' website provides access to plots showing electricity consumption for each device and an option to download the data as a .csv file.

Web Interface

Figure 6 shows an example plot of electricity consumption for one HPWH from the Skycentrics website.

The plot shows that there were two heat pump operation cycles in this monitoring period: approximately 3:30 AM - 5:15 AM and 12:40 PM through the end of the data set. In each case, the heat pump drew about 300 W at the start of the cycle and gradually increased to a little over 500 W by the end of the cycle. This data set matches typical behavior for a HPWH's compressor, which helps validate the quality of this UCM data.

There are occasional timestamps where the CTA-2045 module reported zero electric power in the middle of a heating cycle. These occurrences are likely caused by temporary losses in communication either between the water heater and the CTA-2045 module or between the CTA-2045 module and the cloud.

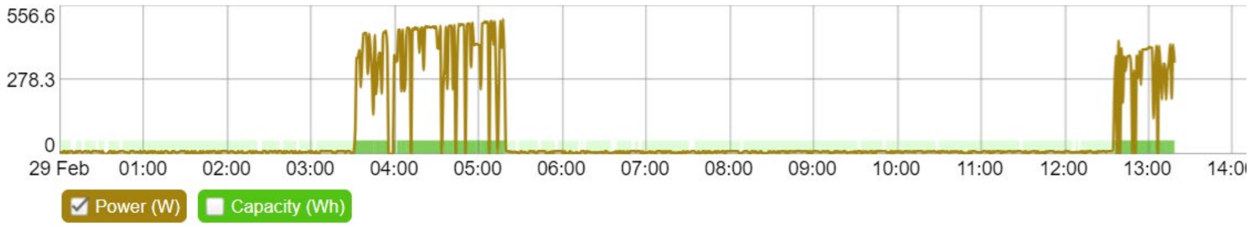

Figure 6. Sample SkyCentrics electricity consumption profile for one HPWH
Data from midnight to a little after 1 PM on February 29, 2024[10]

Data Downloads

The .csv download option provides access to this dataset for offline analysis, presenting one possible option for obtaining baseline electricity consumption for HPWHs included in a fleet-wide load shifting program. However, the interface currently only enables data download or plotting for one HPWH at a time. The time required to generate new plots and download new data sets is a current limitation for scaling up demand response programs using the SkyCentrics platform.

To develop a more efficient data access pathway, the research team asked SkyCentrics to create a batch download process, which enabled a data export for multiple HPWH over a longer time period in one combined file.

SkyCentrics used the batch download process to send baseline data for this project to the research team. Figure 7 shows a comparison of the average weekday power data obtained via the batch download process and the on-site monitoring for a single HPWH at a single site.

---

[10] While the x axis extends a little over 14 hours, the data shown only plot only extends a little over 13 hours.

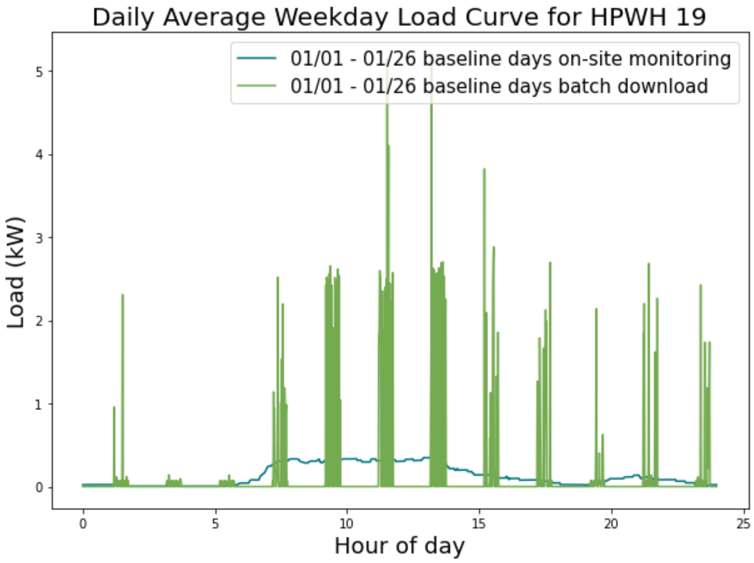

Figure 7. Daily average load curve comparison using on-site monitoring data vs. SkyCentrics batch download data
Data from the HPWH on Site 19, 01/01/2024 through 01/26/2024

We found that the batch download data only outputs the cumulative electricity consumption every 120 minutes. This caused sudden, large, and unrealistic increases in electric power at each timestamp during data processing, as shown in Figure 7.

To improve the usability of the downloaded data, we averaged the increase in cumulative electricity consumption over the time between the prior update and the next update, which is shown in Figure 8. While this alternative approach of processing the batch download data yielded good agreement with the on-site monitoring data, the processed batch download data still had data spikes and low data resolution. Thus, the need to perform data post-processing and the low quality of the processed batch download data pose challenges to widespread deployment.

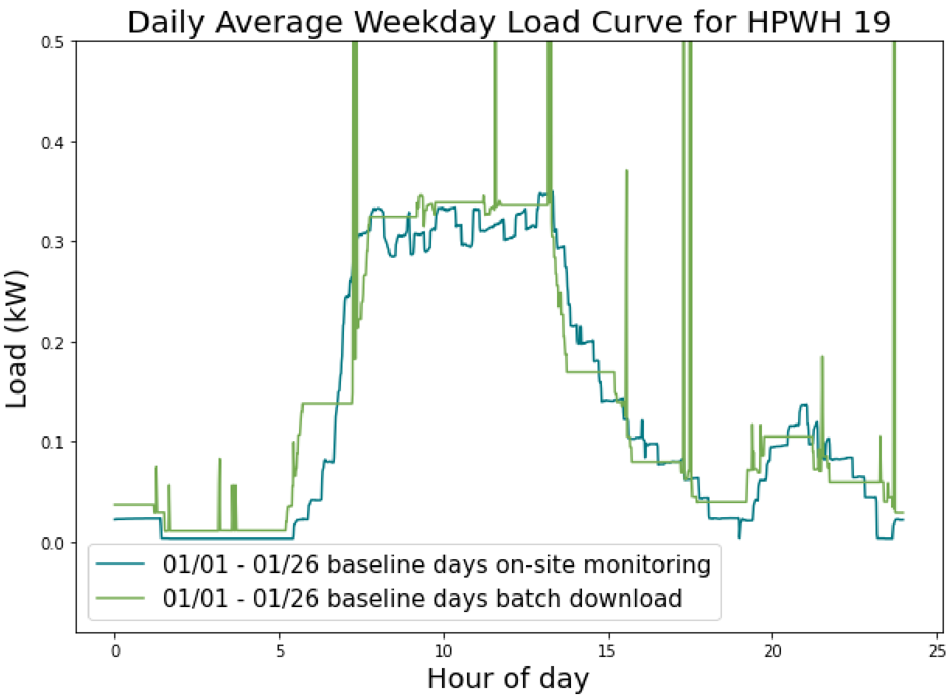

Figure 8. Daily average load curve comparison using on-site monitoring data vs. processed SkyCentrics batch download data
Data from the HPWH on Site 19, 01/01/2024 through 01/26/2024

Load Shifting Performance

The following sections outline the performance of the HPWHs' response to CTA-2045-B signals, and the approximate electricity and cost savings associated with various price schedules.

Response to CTA-2045-B Signals

A HPWH load shifting test was performed for the weekdays between March 4, 2024 and March 8, 2024 for eight HPWHs optimimizing for the CFH WinterHDP price schedule (see Table 2). These eight HPWHs were categorized into four groups based on the daily hot water usage patterns; each group had 1-3 HPWHs. Each group was assigned a different load shifting schedules and received signals at different times of the day. Figure 9 and Figure 10 display HPWH energy and water consumption on an example midweek test day in the middle of the testing period – March 6, 2024 – for two out of the four HPWH groups.

We sent three different control signals to these two groups: Shed, LU, and ALU at different times to the water heaters. The expected behaviors are as follows:

- **Shed**: The heat pump only turns on if the HPWH's tank temperature falls enough to risk a cold-water event.
- **LU**: The heat pump immediately turns on and consumes electricity until the water in the tank is at the setpoint temperature. This could be a short heating cycle if the HPWH is already close to its setpoint temperature.

- **ALU**: The heat pump immediately turns on and consumes electricity until the water in the tank is at an elevated set temperature (exceeding the setpoint). Due to the elevated set temperature, this is expected to be a longer heating cycle than LU.

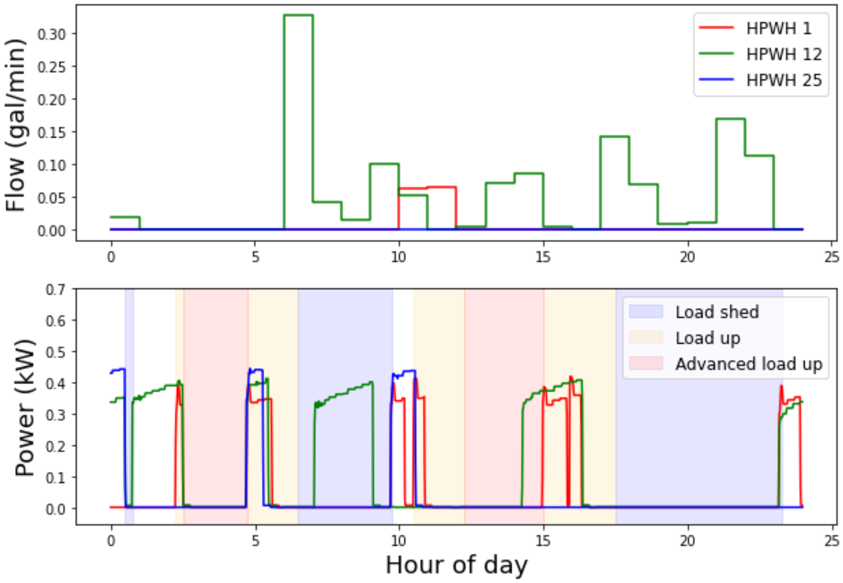

Figure 9. Water consumption and electric power consumption on the example test day 2024-03-06 for HPWHs in the first HPWH group

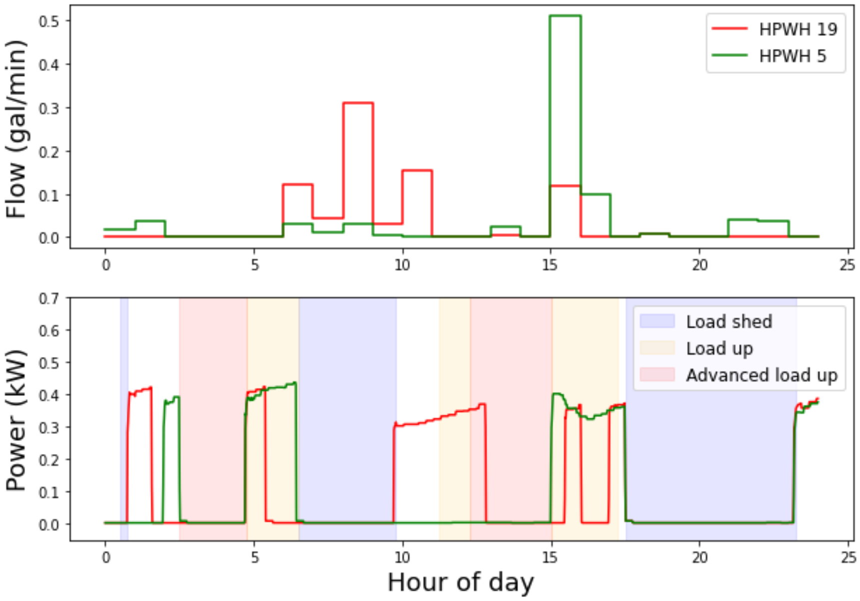

Figure 10. Water consumption and electric power consumption on the example test day 2024-03-06 for HPWHs in the second HPWH group

Figure 9 and Figure 10 both show multiple instances when water heaters immediately turned on at the beginning of LU periods, meaning they responded to LU signals as expected. We then see in multiple instances that the water heaters immediately shut off power at the beginning of load shed periods and immediately turned on at the end of load shed periods. One water heater consumed power during the load shed period, but the occupants consumed nearly 20 gallons of hot water right before the heat pump activated during the load shed period. This was expected for water heaters to avoid cold water events.

However, there are multiple instances where the water heaters immediately deactivated at the beginning of ALU periods and immediately activated at the end of ALU periods. Two water heaters consumed power during the ALU period, but the occupants consumed large volumes of hot water right before power was consumed during the ALU period. This shows that the water heaters didn't respond to ALU signals. It's possible that the HPWHs interpreted them as load shed signals. This is highly undesirable because the water heaters could not be turned on during periods of low electricity prices.

See Appendix A for more information on the additional load shifting potential the ALU signal could unlock if manufacturers ensure their HPWHs respond to signal as requested by the CTA-2045B protocol, i.e., storing additional thermal energy by heating above the setpoint temperature.

Statistical Summary of Load Shifting Performance for All Test Cases

The following sections present findings from each of the field test cases explored in the project. We describe the performance of the controls in response to the different price profiles, highlighting some key performance impacts in the different scenarios. Each price curve was evaluated in terms of the following three key metrics:

- **Peak kWh Reduction**: This metric describes the change in electricity consumption during the peak period. For this project's purposes, the peak period is defined as times when the electricity price is above average. This metric shows the value of the load shifting controls in terms of shifting load away from times when the grid may struggle to meet demand.
- **Solar Peak kWh Increase**: This metric describes the increase in electricity consumption during times of day when photovoltaic systems are producing large amounts of electricity. For this project's purposes, the solar peak is defined as 9 AM to 4 PM. This metric shows the value of load shifting controls in terms of ability to help avoid renewable energy curtailment and to reduce carbon emissions by consuming solar power.
- **Operating Cost Decrease**: This metric describes the change in electricity costs the participants pay. This metric shows the value to the occupants of participating in a load shifting program and describes the economic incentive they would have to participate.

Table 3 presents these three metrics for all five test cases.

Table 3. Key performance metrics for load shifting controls in all price scenarios

Price Schedule	Peak kWh Reduction	Solar Peak kWh Increase	Operating Cost Decrease
CalFlexHub WinterHDP	41%	-14%	30%

CalFlexHub SummerHDP	52%	-13%	46%
SMUD	-5%	12%	-9%
PG&E EV2-A	54%	13%	8%
SCE TOU-D-5-8PM	29%	9%	19%

The test cases using the CalFlexHub prices both showed large decreases in operating cost at 30% and 46%, driven primarily by the high difference between the highest and lowest price. Those differences in the prices caused large shifts in load from the peak period to lower price periods, reducing the peak period electricity consumption by 41% and 52%. Unfortunately, the CalFlexHub prices did not drive the load shifting to solar peak times. Instead, they typically shifted load from the evening peak to the overnight period, and actually reduced the electricity that was consumed during the solar peak by -14% and -13%. Prior simulations, and the simulations shown in this report, strongly imply that this would not be the case if the ALU signal enabled further increases in stored energy during the solar peak period.

Of the utility prices, the PG&E, EV2-A rate yielded the largest reduction in the evening peak period. This effect was driven by the high electricity prices in the evening, a common characteristic of electric vehicle price structures. Those high prices incentivized shifting the load to the overnight period when electricity was cheaper. Through that shift customers were able to reduce their bills by 8% while reducing peak period consumption by 54%. The EV2-A rate also drove an increase of solar-peak electricity of 13% driven by increasing the energy stored in the tank prior to the evening increase in price.

The SCE, TOU-D-5-8PM price schedule yielded the largest decrease in operating costs for the participants, driven by the mid-day decrease in electricity costs. This mid-day decrease created a $0.27/kWh difference between mid-peak and peak, which allowed the participants to consume the lowest cost electricity to pre-charge the tank prior to the peak period. Shifting load from the peak period to that lowest price yielded a 19% reduction in operating costs.

The five subsections below provide a more detailed discussion of each individual test case.

Performance of Load Shifting Controls Responding to CalFlexHub Prices

CalFlexHub WinterHDP

The first test week used CTA-2045A signal schedules minimizing the cost of operating the fleet of HPWHs in response to the *CFH WinterHDP* price schedule. *CFH WinterHDP* is a price schedule which incentivizes load shifting away from the mid-morning and early-evening winter load peaks in California.

This test focused on evaluating 120V HPWH's conformance to the ALU signal. To provide a baseline (case with no ALU signal) the controls for this test assumed that ALU was not available. The options available to the algorithm were the LU and Shed signals. We avoided using the more extreme load shedding signals to a) minimize the chances of participants receiving cold water, and b) assume that grid emergencies are avoided by widespread load shifting controls.

Figure 11 presents the results from the *CFH, WinterHDP, No ALU* test case on a fleet of 10 HPWHs. The blue dashed lines indicate the "unsupervised" load (i.e., no load shifting signal sent, corresponding to the baseline period) and the green solid lines indicate the "supervised" load (i.e., responding to load shift signal during the test week). In both the middle and bottom plots, the red shaded period indicates times when the electricity price is above average, indicating times when the controls should reduce energy consumption.

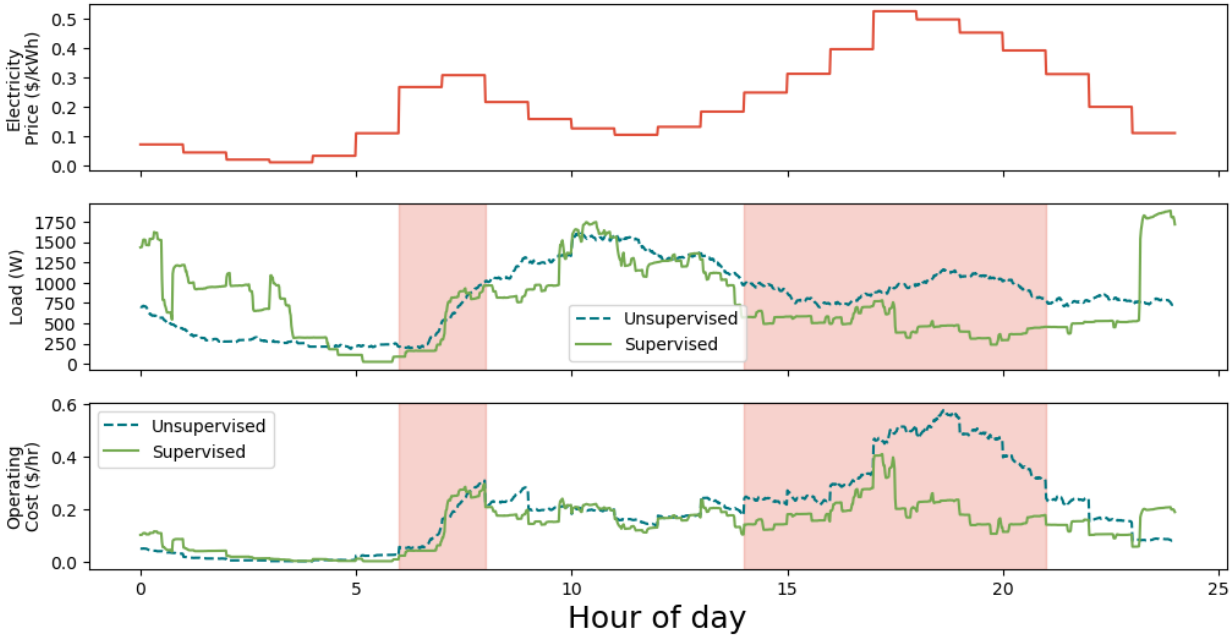

Figure 11. Field results for the *CFH, WinterHDP, No ALU* scenario

Table 4 summarizes the energy and cost impacts of load shifting during the two priority windows. The reduction in demand is especially significant after 5:00 PM, when the electricity prices are at their peak. Successfully shifting electricity consumption from the evening high price period ($0.40/kWh) to the overnight low-price period ($0.05/kWh) yielded an average operating cost reduction of about 30%, as shown in Table 4.

Shifting load away from the morning high price period was not as successful. Despite pre-charging the HPWHs' tanks in response to LU signals, the stored energy was not enough to coast through the morning shower rush without activating the heat pumps. This observation highlights the importance of the ALU signal in CTA-2045B, especially in future electrification scenarios when California has a winter morning peak.

Table 4. *CFH, WinterHDP, No ALU* test case key load shifting statistics

	Peak kWh Reduction	Solar Peak kWh Increase	Operating Cost Decrease
CFF, WinterHDP, No ALU Impacts	42%	-14%*	30%

**Negative number indicates that solar peak time energy usage was decreased (not increased)*

CFH SummerHDP

Figure 17 summarizes the results of the *CFH, SummerHDP* price schedule on a fleet of 10 HPWHs, which was tested after the *WinterHDP* test case. Because the *SummerHDP* price profile does not include a high price period in the morning, the CTA-2045 signal schedules for this case did not feature morning LU periods[11].

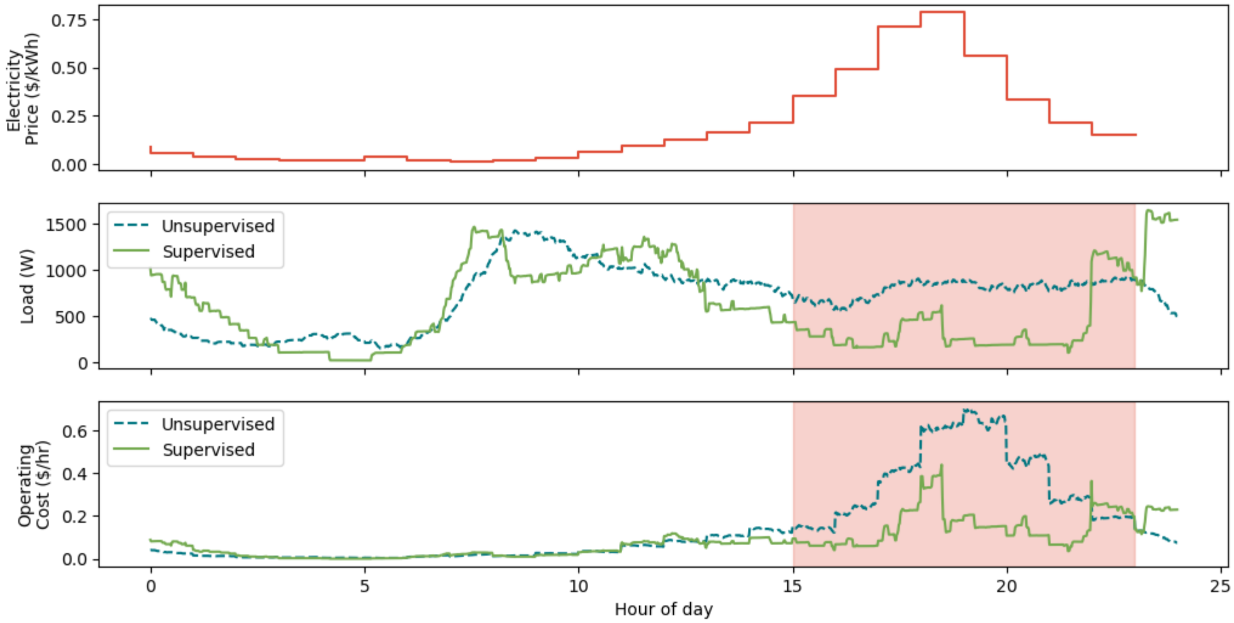

Figure 12. Field results for the *CalFlexHub, SummerHDP, No ALU* scenario

Similarly to the *WinterHDP* scenario, the controls for the *SummerHDP* test case successfully shifted load out of the evening high price period (average price of $0.21/kWh) and into the overnight period (average price of $0.13/kWh).

Table 5. *CFH, SummerHDP, No ALU* test case key load shifting statistics

	Peak kWh Reduction	Solar Peak kWh Increase	Operating Cost Decrease
CFH, SummerHDP, No ALU Impact	52%	-13%*	46%

**Negative number indicates that solar peak time energy usage was reduced (not increased)*

This field test displayed a 2-stage snapback period after the conclusion of the high price period. The HPWHs quickly activated their heat pumps as the electricity price decreased and the load shifting controls released them from Shed mode. This yielded spikes in electricity consumption at 10:00 and 11:00 PM which suggest that scaled programs should consider a staggering of the

---

[11] California's utility grid does not currently have a summer morning peaking problem, so there is not a financial incentive to pre-charge the tank before the morning shower period.

controls to avoid a sudden spike. One possible approach would be releasing a few HPWHs from Shed mode every few minutes.

After turning on at ~10-11 PM, the heat pumps recharged until ~2-3 AM. This meant the HPWHs consumed less electricity from ~4-6 AM than the baseline (unsupervised) case, so the tanks contained less thermal energy. This led to the HPWHs activating their heat pumps more rapidly and consuming more electricity during the morning shower period.

The controls successfully shifted the load away from the evening high price period, and towards the overnight period. The test case showed a reduction in electricity consumption during the high price period of 52%, resulting in an operating cost reduction of 46%. However, with this price schedule and ALU the control shifted load to the overnight period instead of the mid-day period, resulting in a 13% decrease in mid-day solar power absorbed.

Performance of Load Shifting Controls Responding to Investor Owned Utilities (IOU) Time of Use (TOU) Prices

The final weeks of testing evaluated the load shifting performance of LBNL's controls for 120V HPWHs in response to the three current IOU TOU rates. Since the three rate structures have different features, these three experiments provide some insight into how different rate structures impact load shifting performance.

PG&E EV2-A

Figure 13 presents the impact of load shifting controls under PG&E's EV2-A price structure. This price structure was selected to represent a future scenario: as more Californians purchase electric vehicles more customers will use electric vehicle rate plans. These rate plans feature highly elevated electricity prices from 3PM to midnight, and reduced prices the rest of the day, encouraging customers to shift their load away from evening peak periods.

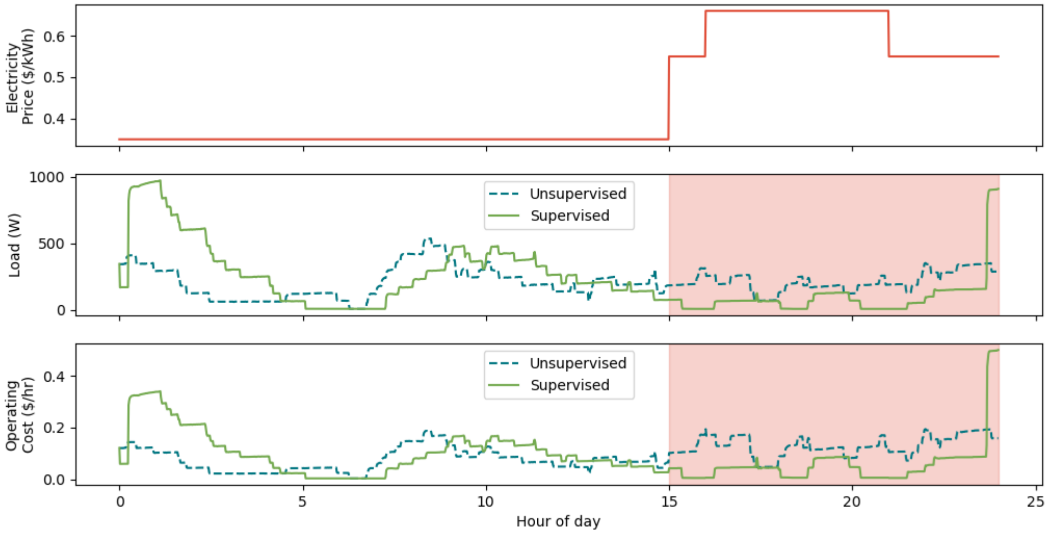

Figure 13. Field results for the *PG&E EV2-A, Summer, No ALU* test case

The middle plot shows the load for the fleet of four HPWHs tested with PG&E's EV2-A rate. The controls successfully shifted load away from the high price period. The results with the load shifting controls consumed more than twice as much electricity between midnight and 5AM as compared to the baseline (unsupervised) case, due to reheating the water in the tank after the high price period. However, since the price of electricity is reduced overnight, the increased load only caused small increases in the operating cost.

Table 6. *PG&E EV2-A, Summer, No ALU* test case key load shifting statistics

	Peak kWh Reduction	Solar Peak kWh Increase	Operating Cost Decrease
EV Tariff Impacts	54%	13%*	8%**

**Negative number indicates that solar peak time energy usage was reduced (not increased)*

***Post-processing calculations estimate that improvements to the control logic would reduce the average cost of electricity to $0.38/kWh and increase the operating cost reduction to 10%*

As shown on the right-hand side of the middle plot, all the heat pumps activated at 11:30 PM instead of waiting until the electricity price reduced at midnight. This simultaneous activation is driven by a flaw in the control logic. Since the controls were originally developed to work with gradual changes in prices, akin to the CalFlexHub price schedules, they did not correctly respond to the less frequent & larger price changes in the EV2-A rate. Since the testing uncovered this issue, LBNL has resolved a bug that incorrectly interpolated prices causing the controls to release from Shed early, and this issue should not occur in future deployments.

SCE TOU-D-5-8PM (Winter)

Figure 14 shows the same results for the three HPWHs responding to SCE's TOU-D-5-8PM, Winter price schedule. This price schedule was selected to evaluate the impacts of having a midday low price period in the price schedule. This low price in the middle of the day should further encourage loading up the HPWHs prior to the high price period, both making the load shifting more effective and increasing the absorption of mid-day solar power.

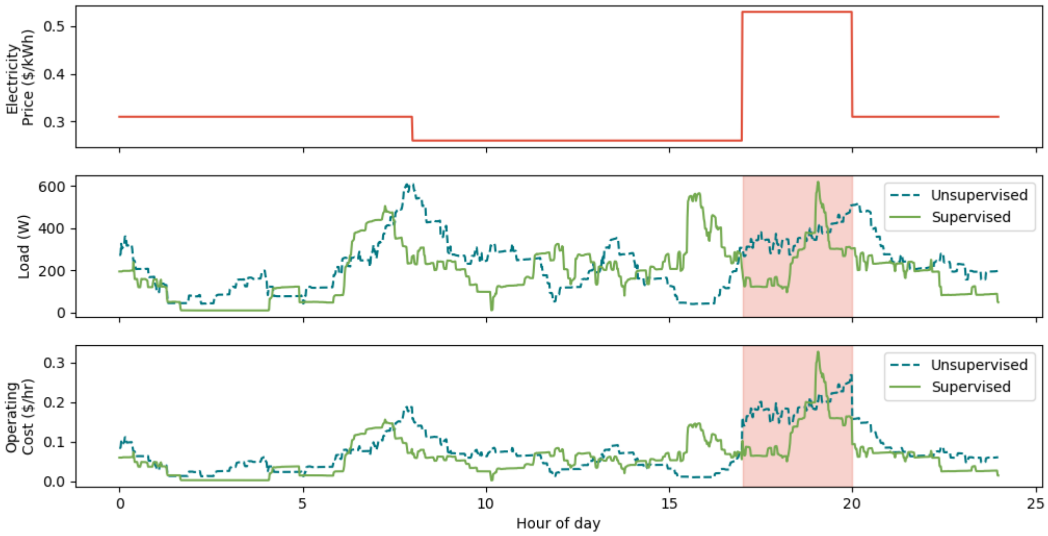

Figure 14. Field results for the *SCE TOU-D-5-8PM, Winter, No ALU* test case

The reduced mid-day price for this test case caused an increase in load prior to the peak period, showing an increase in average demand from ~50W during the baseline period to ~525W during the test case. This increase in demand during a load-up period both increased the potential to absorb mid-day solar power by 9.4% and ensured the storage tanks were charged entering the high price period. The increase in energy stored in the tank allowed better ability to shift load away from the peak period.

Table 7. *SCE TOU-D-5-8PM, Winter, No ALU* test case key load shifting statistics

	Peak kWh Reduction	Solar Peak kWh Increase	Operating Cost Decrease (%)
Low Mid-Day Tariff Impacts	29%	9%	19%

The load plot does show a spike in electricity consumption during the peak period, prior to the controls sending a Shed signal at 7PM. The control algorithm did not find that a Shed period through the entire high price period yielded optimal results. Instead, it selected a 1 hour Shed only at the end of the high price period. As can be seen in Figure 14, during the test period the users consumed enough hot water to drive peak loads up to 600 W during the high price period. This is a surprising result which may be caused by efforts to avoid cold water events. Further evaluation should explore the best way to address these situations. Two options include: 1) Accepting the output signal schedule as is, or 2) Post-process the output to Shed through the entire high price period, yielding better load shifting at the cost of increasing the risk of cold water events.

The bottom plot shows the operating cost of the fleet both with and without load shifting controls. Shifting 29% of the load from the on-peak, $0.52/kWh, prices to the mid-day, $0.25/kWh, prices reduced the cost of operating the HPWHs by 19%.

SMUD (Summer)

Figure 15 shows the same plots for the SMUD, Summer test scenario with a fleet of three HPWHs. In this case the electricity price was lowest overnight with a mild increase at mid-day prior to a large peak price during the 5-8PM evening peak window. The price returned to the mid-peak price at 8 PM. The difference between the peak period price and the mid-day mid-peak price should incentivize shifting load from the peak period to the mid-day or evening period; however, this effect is likely to be less dramatic than in the SCE, TOU-D-5-8PM case because the mid-day price was the lowest in that case.

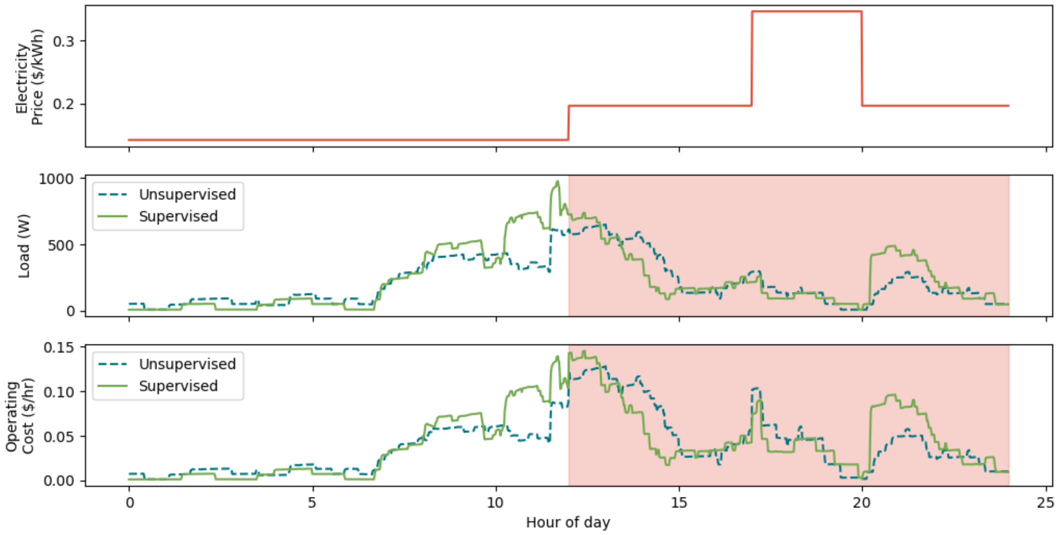

Figure 15. Field results for the *SMUD, Summer, No ALU* test case

The load shifting controls successfully initiated a load-up period prior to the noon increase in electricity price, which decreased the electricity consumption during the 2-3 PM window. However, the electricity consumption in the cases with and without load shifting control are nearly identical from 3 PM to 8 PM, implying that the hot water draws were large enough to consume the energy stored in the tank and require additional heating during the elevated price period. The load shifting controls included a Shed signal during the evening peak period, which reduced the load during the evening peak by a marginal amount. As the price reduced from the peak to the mid-peak, the devices were released from Shed mode causing a rebound which increased the electricity consumption during the mid-peak period.

These impacts were exacerbated, and perhaps caused, by an increase in hot water consumption during the test period. In the baseline period the homes consumed an average of 34.8 gal/day to hot water, which increased to 39.4 gal/day (+15%) during the test period. The consumption during the high price period increased from 4.8 gal/day to 6.0 gal/day (+26%). It could also be caused by the control algorithm not looking far enough ahead in the price; it's possible that a longer forecast horizon would have enabled delaying the rebound until just after midnight which would have shifted a portion of the load to a lower price time of day.

Driven by these impacts, the period with the load shifting controls showed an 9% increase in operating costs. Taking the increase in consumption into account, the load shifting controls showed a reduction in operating costs from $0.027/gal to $0.026/gal, even though there was a 26% increase in consumption during the most expensive part of the day.

	Peak kWh Reduction	Solar Peak kWh Increase	Operating Cost Decrease
Graduated Price Tariff Impacts	-5%*	13%	-9%*

**Post-processing calculations identified that the occupants consumed 15% more hot water and 26% more in the peak period during the test period than the baseline period. This behavioral change drove a 5% increase in peak kWh consumption and a 9% increase in HPWH operating costs. These are shown as negative decreases for consistency with other tables. Post-processing estimates that without the increase in consumption the peak kWh reduction would be 21% and the operating cost decrease would be 18%.*

Table 8. *SMUD, Summer, No ALU* test case key load shifting statistics

CONCLUSIONS

The primary aim of this project was to assess the cost and peak load shaving efficacy of 120V HPWH load shifting. To accomplish this task, the study team developed and deployed custom load shifting controls optimized for two CFH highly dynamic price schedules and three California utility price schedules. The study team effectively shifted at least 29% of energy use, on average, away from the evening peak period for four out of five price schedules. This power curtailment was achieved despite the nonconformance of the ALU signal. With ALU, the HPWHs could have stored additional thermal energy, and likely shifted even more energy away from high price periods.

The following sections provide more detail on costs versus benefits of UCM load shifting, the need for ALU conformance, and recommendations for scaling up a similar HPWH demand response program.

Upfront Costs

The primary downside to CTA-2045 UCMs at present is cost. At $150-$209 each, it can be difficult to justify implementing this technology. According to previous literature on HPWH load shifting, 240V HPWHs can consistently shift between 200 and 750 watt-hours (Wh) from high-cost hours to low-cost hours. 120-volt HPWHs in this study fall on the lower end of that range due to their smaller or non-existent resistance elements. Therefore, it takes hundreds of load shifting events to recover hardware costs based on current time-of-use price structures.

However, as the need for behind-the-meter demand response grows, the value of load shifting services will also increase. In tandem, UCM prices will likely fall as manufacturers receive larger orders and increase economies of volume.

Advanced Load Up

This project uncovered that the 120V HPWHs tested in this study do not respond correctly to ALU signals sent by SkyCentrics' platform. The cause of this disconnection is not currently clear. Either the HPWHs are not compliant with the CTA-2045B protocol or there is a communication issue between SkyCentrics and the HPWHs.

To achieve the full potential of load shifting controls HPWHs, future efforts should focus on increasing adoption of the ALU signal. This could be pursued through a combination of ideas like advocating for adoption of CTA-2045B, providing larger program participation incentives for CTA-2045B compliant products, or changing regulations to require adoption of CTA-2045B.

Scalability

At the start of this project, only one manufacturer had 120V HPWH products available on the market; thus, the study team focused on products from that manufacturer. Since then, one other major manufacturer has released a market-ready product, and others continue to work on prototypes. Scaling these findings to large-scale programs will require ensuring the controls work well for all manufacturers. Further laboratory or field tests enabling these controls to support those products are recommended.

Procurement and installation of CTA-2045 communication modules are a challenge which could limit adoption of widespread programs. The modules purchased for this project cost $150 each, on top of SkyCentrics data/administration fees, which is a significant cost compared to the load shifted by a HPWH. During this project, we received utility feedback that available modules are too expensive, and that the industry needs less expensive solutions.

Further, since the modules are not pre-installed in the HPWH, it is necessary for building occupants to install and configure the modules themselves. Since SkyCentrics did not clearly communicate the installation and commissioning process, the research team created a separate instruction document to share with participants. Despite this step, participants still struggled with the installation process and the project team had to spend extra time on troubleshooting emails and video calls.

Deploying these controls in large programs will require resolving these two issues with the modules. Developments should focus on incorporating low-cost communication devices directly in the HPWH product. Incorporating CTA-2045 modules directly in the product would both remove the installation challenges and reduce the cost of modules through economies of volume.

Technology Transfer

While this project demonstrated the potential benefits of price- and load-responsive controls for 120V HPWHs, the techniques developed in this project are currently too cumbersome to deploy at scale. However, the grouping method developed in this project and signal schedule generation tool could be useful for future utility programs. LBNL, WatterSaver, and SmartShift have started direct conversations identifying the best way to convert these methods into user-friendly tools to facilitate adoption by those programs. The WatterSaver team expects to try a version of the tools in Q3 of FY24.

LBNL is currently funded by the US Department of Energy Stor4Build[12] project to develop a user-friendly tool that generates groups of HPWHs with similar baseline electricity consumption patterns. When published, this tool will enable users to identify HPWHs that should receive similar load shifting signals and customize controls accordingly. The target is to publish an alpha version of the tool by the end of 2024. Utility programs – e.g., PG&E's WatterSaver or SCE's Smart Shift – testing the tool and providing feedback would help improve the tool and the program's ability to adopt these techniques.

Simultaneously, the analytical framework to develop load shifting controls prior to deployment needs to be improved prior to widespread adoption. The toolchain currently requires Python programming skills, knowledge of the toolchain environment, and can take several days to simulate. This tool must be simplified, streamlined, and sped up to be leveraged by programs.

These two tasks could be facilitated by strong collaboration between the research team and the utility program teams. Further support of the LBNL-led HPWH Load Shifting Working Group would facilitate these conversations and provide the valuable two-way information flow to support improvement and adoption of the tools.

---

[12] See the Department of Energy's Stor4Build website page for more information: https://www.energy.gov/eere/buildings/stor4build

APPENDIX A

The following three sections provide additional information on: LBNL's load shifting algorithm iteration process; SkyCentrics' demand response scheduling tools; and, lastly, the *Advanced Load Up* signal and the additional load shifting potential for HPWHs if they conform to the signal.

Load Shifting Algorithm Iteration Process

Creating groups of HPWHs with similar behavior enables researchers to customize load shifting controls to those behaviors and is one of the keys to improving load shifting performance. Figure 16 shows a visual representation of both the original grouping method and the new, improved method. The original method grouped HPWHs by average kWh consumed each day. This method effectively separated the HPWHs which used backup resistance elements from those which only used the heat pump, as can be seen by comparing Figure 16 Subplot 3 to Subplots 1 and 2. However, displayed in those three subplots, the method did not effectively separate HPWHs by the time of day they consumed electricity.

The new method identifies similarity between HPWH electricity consumption profiles by calculating the root mean square error (RMSE) between different HPWH profiles. The RMSE is sensitive to the difference at each timestep, meaning that electricity consumption at different times of day creates significant differences in RMSE. Groups of HPWHs are then created by identifying the HPWHs with low RMSE values. The impacts can be seen in subplots 4-6. Subplot 6 shows that, like the old method, HPWHs that use the resistance elements are still separated from those that use only the heat pump. Subplots 4, 5, and 6 all show that the groups feature very tightly clustered electricity consumption.

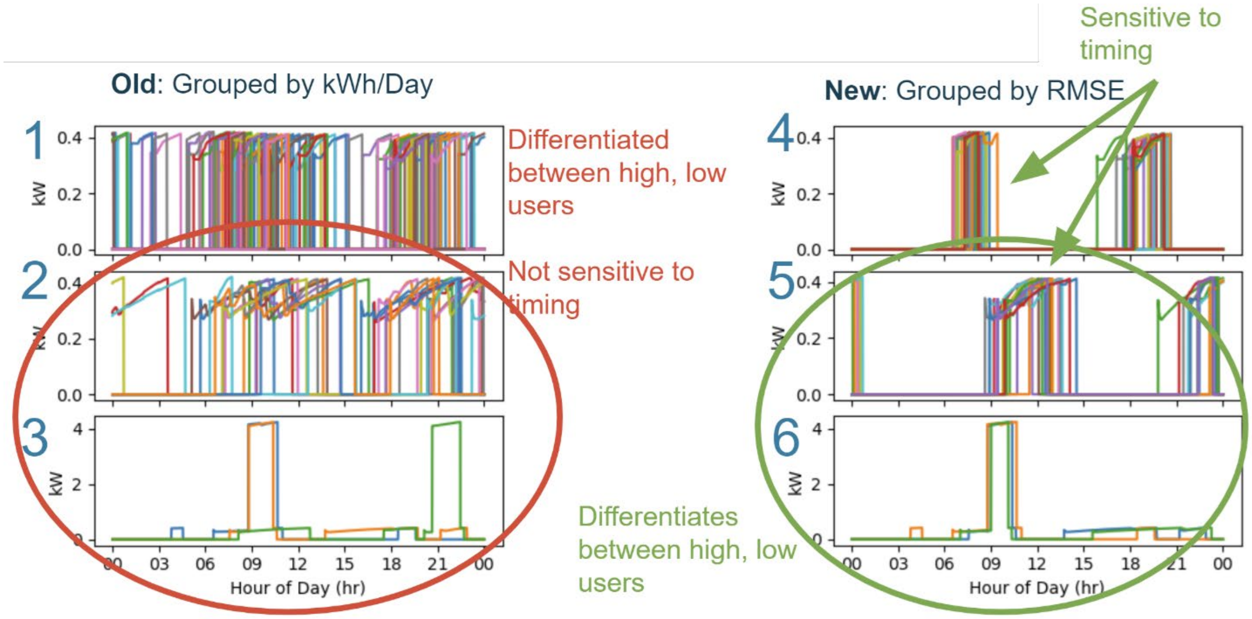

Figure 16. Comparison of old and new grouping methods

Deploying the load shifting algorithm in the field then required a method for adding new, field-deployed HPWHs to the groups of simulated HPWHs. For this project, LBNL developed a method that does so by comparing the typical load curve of a field-deployed HPWH to the load curves in each group. The method is shown visually in Figure 17 and uses the following steps:

1. **Identify typical baseline load curve**: This step calculated the average electricity consumption, expressed in kW, or a specific HPWH for each hour of the day across the data set. This methodology created a load curve representing the typical electricity consumption of a HPWH on any given day.
2. **Calculate the average RMSE for each group:** The next step is to create an error metric expressing how different the consumption of that HPWH is from the HPWHs in each group. To do so, the method calculates the RMSE between the field HPWH and each HPWH in each group. The RMSE for each group is then averaged, providing an average RMSE value comparing the field HPWH to each group.
3. **Select the lowest RMSE group**: Since RMSE quantifies the difference between two data sets, the group with the lowest average RMSE contains HPWHs with load curves that are most similar to the field HPWH. The load shifting signal schedule for that group will yield the best results for the field HPWH, and the field HPWH is then sent the load shifting schedule for that group.

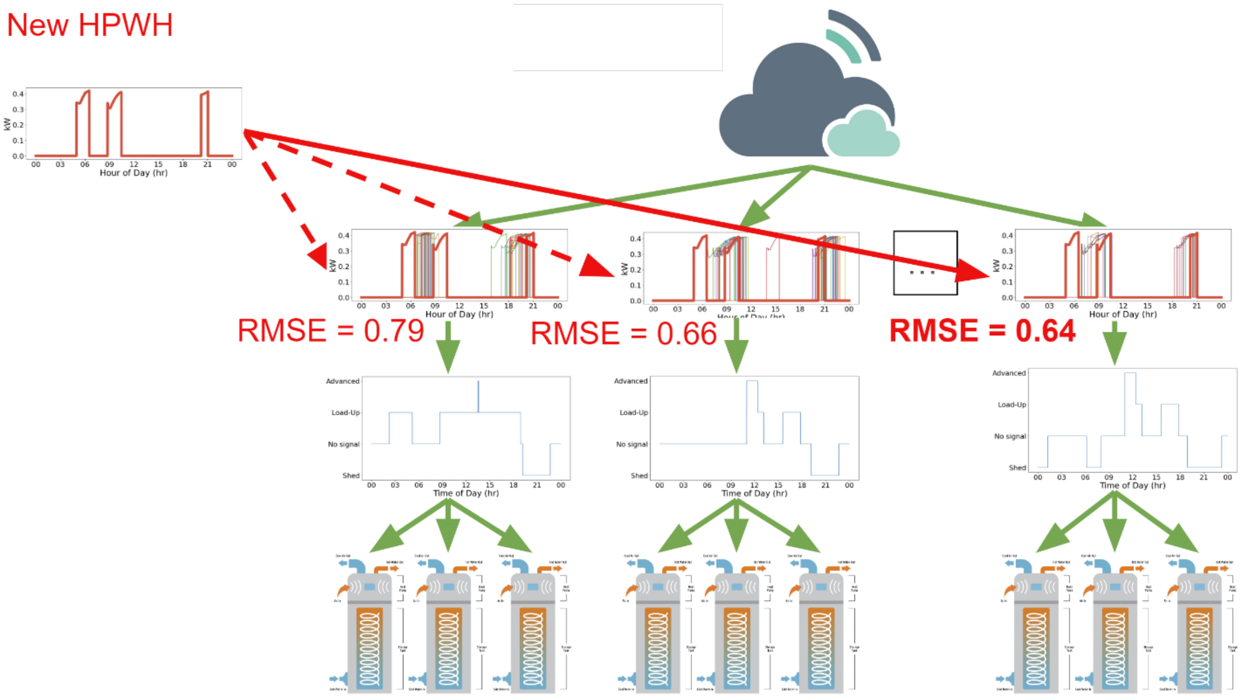

Figure 17. Conceptual schematic showing the method of identifying the signal schedule for a field HPWH

Load Shifting Scheduling

LBNL used the new grouping methods, as described in the Load Shifting Algorithm Iteration Process section, and their tool chain for developing load shifting control strategies to develop new schedules for the 120V HPWHs.

SkyCentrics' web interface provides the capabilities needed to implement price- and load-responsive load shifting controls for HPWHs. Specifically:

- The Device Trees feature enables users to place CTA-2045 modules (effectively HPWHs) into groups and send identical signals to each HPWH in the group.
- The DR Scheduler enables sending the key CTA-2045 signals to the various groups of HPWHs.

Figure 18 presents a screenshot showing the capabilities available for scheduling load shifting events in the DR Scheduler on SkyCentrics' website. The "Create Event" box shows the options to create a new load shifting event. This interface allows the user to 1) specify a name, 2) state that date/time of the event, 3) specify if/when the event should repeat, 4) the type of event (e.g. Shed, LU, etc.), and 5) the group of devices receiving this particular event. Creating a new event using this interface will add a new bar to the calendar and send the load shifting signals to the HPWH at the scheduled time.

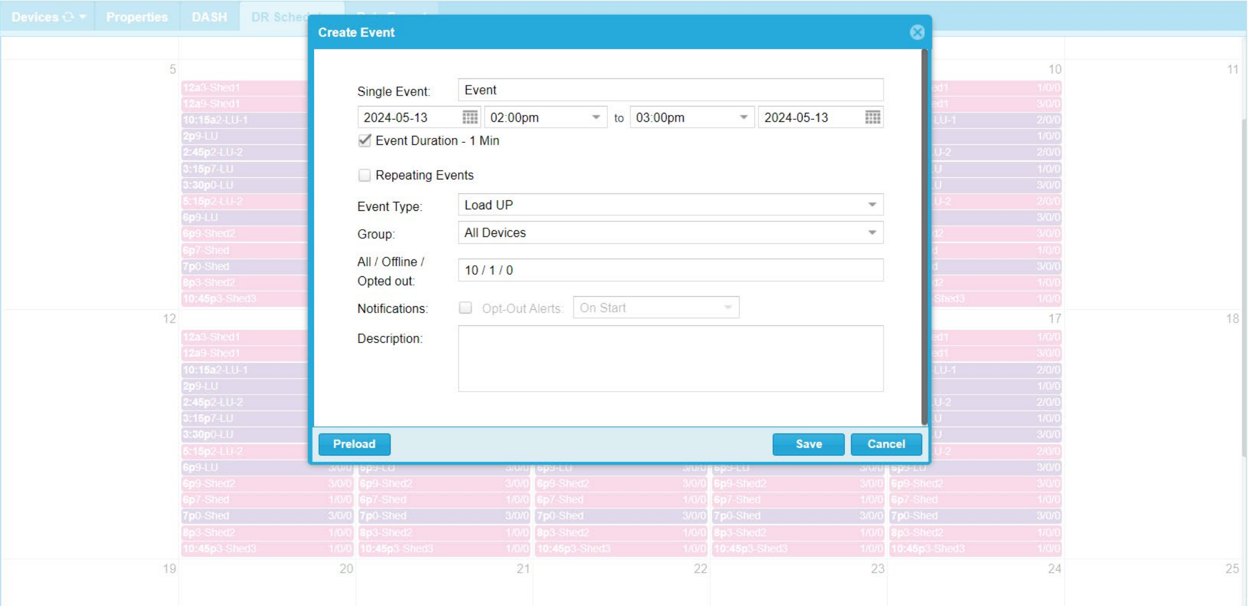

Single Event: Event
2024-05-13 02:00pm to 03:00pm 2024-05-13
☑ Event Duration - 1 Min
☐ Repeating Events
Event Type: Load UP
Group: All Devices
All / Offline / Opted out: 10 / 1 / 0
Notifications: ☐ Opt-Out Alerts: On Start
Description:

Figure 18. Screenshot showing the tools for inputting load shifting schedule on SkyCentrics' website

Figure 19 displays the Devices Tree, which shows how the different CTA-2045 modules are organized. Each folder represents a different group, and each flame icon represents a CTA-2045 module. Dragging a module into a specific folder creates groups of multiple devices, as shown in the screenshot. For example, the group "Zero" has CTA-2045 modules "F14CDE" and "D3375A," and "D33742." This organization structure enables sorting the HPWHs into groups and treating them similarly.

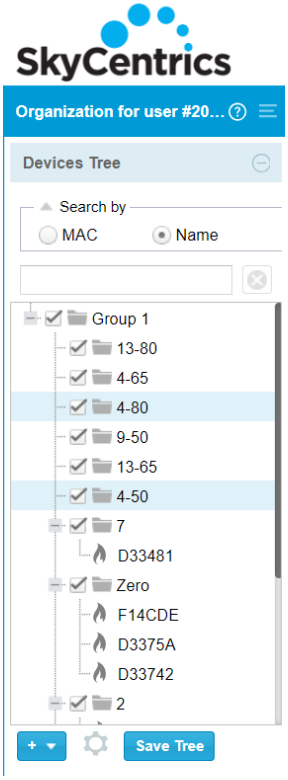

Figure 19. Screenshot showing the SkyCentrics Device Tree

The DR Schedular calendar displayed in Figure 20 shows the currently scheduled load shifting signals. Each colored bar within a given day represents a different signal. The text displayed on each bar shows both the timing of an event and the user-specified name. For instance, the top bar on May 6, 2024, reads “12a3-Shed1” indicating that the signal will be sent at 12 AM (midnight) and has been named “3-Shed1” by the user. Different colors represent different signals – pink bars indicate Shed events, purple bars indicate LU events, and yellow bars indicate ALU (no ALU signals are shown in Figure 20).

May 2024

Sun	Mon	Tue	Wed	Thu	Fri	Sat
28	29	30	1	2	3	4
5	6 12a3-Shed1 1/0/0 12a9-Shed1 3/0/0 10:15a2-LU-1 2/0/0 2p9-LU 1/0/0 2:45p2-LU-2 2/0/0 3:15p7-LU 1/0/0 3:30p0-LU 3/0/0 5:15p2-LU-2 2/0/0 6p9-LU 3/0/0 6p9-Shed2 3/0/0 6p7-Shed 1/0/0 7p0-Shed 3/0/0 8p3-Shed2 1/0/0 10:45p3-Shed3 1/0/0	7 12a3-Shed1 1/0/0 12a9-Shed1 3/0/0 10:15a2-LU-1 2/0/0 2p9-LU 1/0/0 2:45p2-LU-2 2/0/0 3:15p7-LU 1/0/0 3:30p0-LU 3/0/0 5:15p2-LU-2 2/0/0 6p9-LU 3/0/0 6p9-Shed2 3/0/0 6p7-Shed 1/0/0 7p0-Shed 3/0/0 8p3-Shed2 1/0/0 10:45p3-Shed3 1/0/0	8 12a3-Shed1 1/0/0 12a9-Shed1 3/0/0 10:15a2-LU-1 2/0/0 2p9-LU 1/0/0 2:45p2-LU-2 2/0/0 3:15p7-LU 1/0/0 3:30p0-LU 3/0/0 5:15p2-LU-2 2/0/0 6p9-LU 3/0/0 6p9-Shed2 3/0/0 6p7-Shed 1/0/0 7p0-Shed 3/0/0 8p3-Shed2 1/0/0 10:45p3-Shed3 1/0/0	9 12a3-Shed1 1/0/0 12a9-Shed1 3/0/0 10:15a2-LU-1 2/0/0 2p9-LU 1/0/0 2:45p2-LU-2 2/0/0 3:15p7-LU 1/0/0 3:30p0-LU 3/0/0 5:15p2-LU-2 2/0/0 6p9-LU 3/0/0 6p9-Shed2 3/0/0 6p7-Shed 1/0/0 7p0-Shed 3/0/0 8p3-Shed2 1/0/0 10:45p3-Shed3 1/0/0	10 12a3-Shed1 1/0/0 12a9-Shed1 3/0/0 10:15a2-LU-1 2/0/0 2p9-LU 1/0/0 2:45p2-LU-2 2/0/0 3:15p7-LU 1/0/0 3:30p0-LU 3/0/0 5:15p2-LU-2 2/0/0 6p9-LU 3/0/0 6p9-Shed2 3/0/0 6p7-Shed 1/0/0 7p0-Shed 3/0/0 8p3-Shed2 1/0/0 10:45p3-Shed3 1/0/0	11
12	13 12a3-Shed1 1/0/0 12a9-Shed1 3/0/0 10:15a2-LU-1 2/0/0 2p9-LU 1/0/0 2:45p2-LU-2 2/0/0 3:15p7-LU 1/0/0 3:30p0-LU 3/0/0 5:15p2-LU-2 2/0/0 6p9-LU 3/0/0 6p9-Shed2 3/0/0 6p7-Shed 1/0/0 7p0-Shed 3/0/0 8p3-Shed2 1/0/0 10:45p3-Shed3 1/0/0	14 12a3-Shed1 1/0/0 12a9-Shed1 3/0/0 10:15a2-LU-1 2/0/0 2p9-LU 1/0/0 2:45p2-LU-2 2/0/0 3:15p7-LU 1/0/0 3:30p0-LU 3/0/0 5:15p2-LU-2 2/0/0 6p9-LU 3/0/0 6p9-Shed2 3/0/0 6p7-Shed 1/0/0 7p0-Shed 3/0/0 8p3-Shed2 1/0/0 10:45p3-Shed3 1/0/0	15 12a3-Shed1 1/0/0 12a9-Shed1 3/0/0 10:15a2-LU-1 2/0/0 2p9-LU 1/0/0 2:45p2-LU-2 2/0/0 3:15p7-LU 1/0/0 3:30p0-LU 3/0/0 5:15p2-LU-2 2/0/0 6p9-LU 3/0/0 6p9-Shed2 3/0/0 6p7-Shed 1/0/0 7p0-Shed 3/0/0 8p3-Shed2 1/0/0 10:45p3-Shed3 1/0/0	16 12a3-Shed1 1/0/0 12a9-Shed1 3/0/0 10:15a2-LU-1 2/0/0 2p9-LU 1/0/0 2:45p2-LU-2 2/0/0 3:15p7-LU 1/0/0 3:30p0-LU 3/0/0 5:15p2-LU-2 2/0/0 6p9-LU 3/0/0 6p9-Shed2 3/0/0 6p7-Shed 1/0/0 7p0-Shed 3/0/0 8p3-Shed2 1/0/0 10:45p3-Shed3 1/0/0	17 12a3-Shed1 1/0/0 12a9-Shed1 3/0/0 10:15a2-LU-1 2/0/0 2p9-LU 1/0/0 2:45p2-LU-2 2/0/0 3:15p7-LU 1/0/0 3:30p0-LU 3/0/0 5:15p2-LU-2 2/0/0 6p9-LU 3/0/0 6p9-Shed2 3/0/0 6p7-Shed 1/0/0 7p0-Shed 3/0/0 8p3-Shed2 1/0/0 10:45p3-Shed3 1/0/0	18

Figure 20. Screenshot showing scheduled demand response events in SkyCentrics

The research team used this interface and these capabilities to send the load shifting controls to the HPWHs. While this interface was appropriate for a research project of this size, scaling it up to a widespread program would be challenging. Each signal had to be programmed into the interface manually, which is time consuming and creates opportunities for manual errors and incorrectly sent signals. Using the interface to program signals for the many HPWHs in a large program would require significant labor time and lead to inappropriately sent signals. Use and/or development of a direct API would enable connecting the computer program generating the signal schedules directly to the signals sent, both speeding up the process and reducing the risks of entry mistakes.

Advanced Load Up Simulation Analysis

To determine the impacts of the *Advanced Load-Up* signal not being supported by all manufacturers, the research team performed simulations showing the potential difference in performance. These simulations assumed that all HPWHs had 80-gallon storage tanks, a set temperature of 125 °F, and that the ALU signal temporarily increases the set temperature to 140°F. The simulations assumed 148 HPWHs, each with a different monitored hot water draw profile. LBNL used the algorithm to determine cost-minimizing load shifting control schedules and the resulting impacts for two different cases: one case assumptions that ALU is available, and the second case assumes that it is not. These simulations utilized the *CalFlexHub SummerHDP* price schedule.

Figure 21 and Figure 22 show the simulation results both without (Figure 21) and with (Figure 22) the ALU signal. The case without ALU shows only a small increase in consumption during the late morning/early afternoon period. This is driven by the occupants taking morning showers. Most heat pumps were already operating to recover from the showers, and the LU signal was only able to drive a little more energy storage. Shed signals sent in the evening had a larger impact, decreasing the electric demand between 5 and 9 PM. This impact can be seen especially clearly in the cost subplot, with the supervised case showing a noticeable decrease in fleet operating cost. Over the 24-hour day, the load shifting controls reduced the operating cost by 32%, primarily by reducing the peak period consumption by 29%. Mid-day solar consumption increased by 14%.

The case with ALU shows larger differences. Driven by the very low morning electricity prices in the *CalFlexHub SummerHDP* price schedule the control algorithm specified to use *Advanced Load-Up* in the 7 AM to noon window. This created a large increase in morning electricity consumption, almost double the baseline case, as the heat pumps operated for longer periods, causing more overlap with other heat pumps. This longer heat pump operation stored more energy in the tank, which enabled better load shedding performance in the evening high price period. The result is better performance on all metrics. Including *Advanced Load-Up* enabled decreasing the operating cost by 53%, decreasing high price period consumption by 60%, and increasing mid-day solar consumption by 65%.

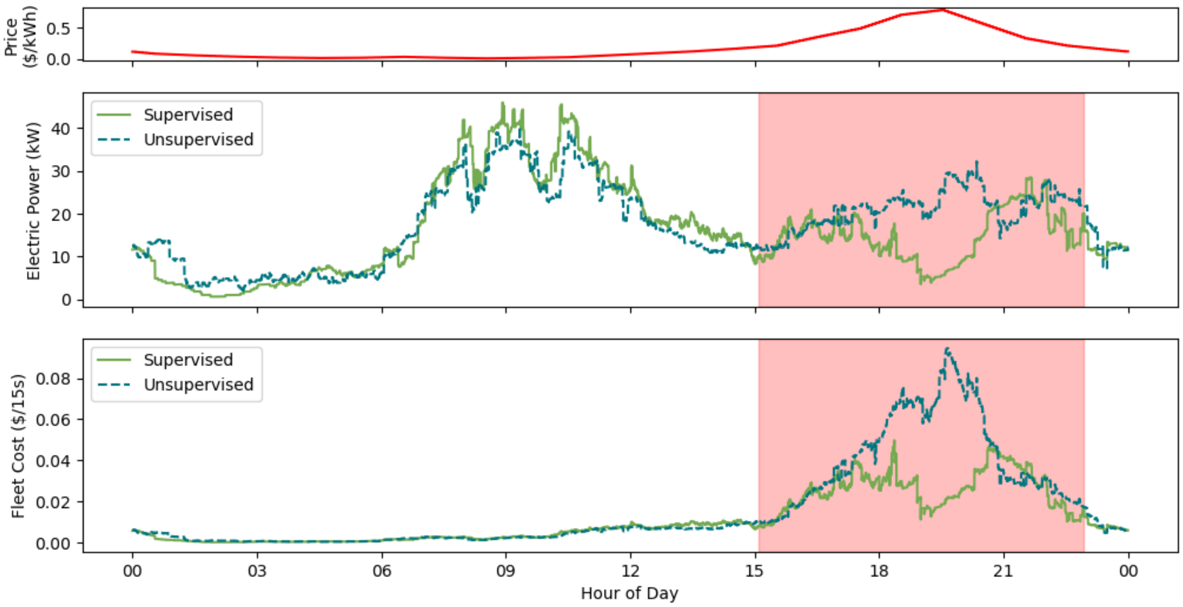

Figure 21. Simulations results assuming no ALU signal

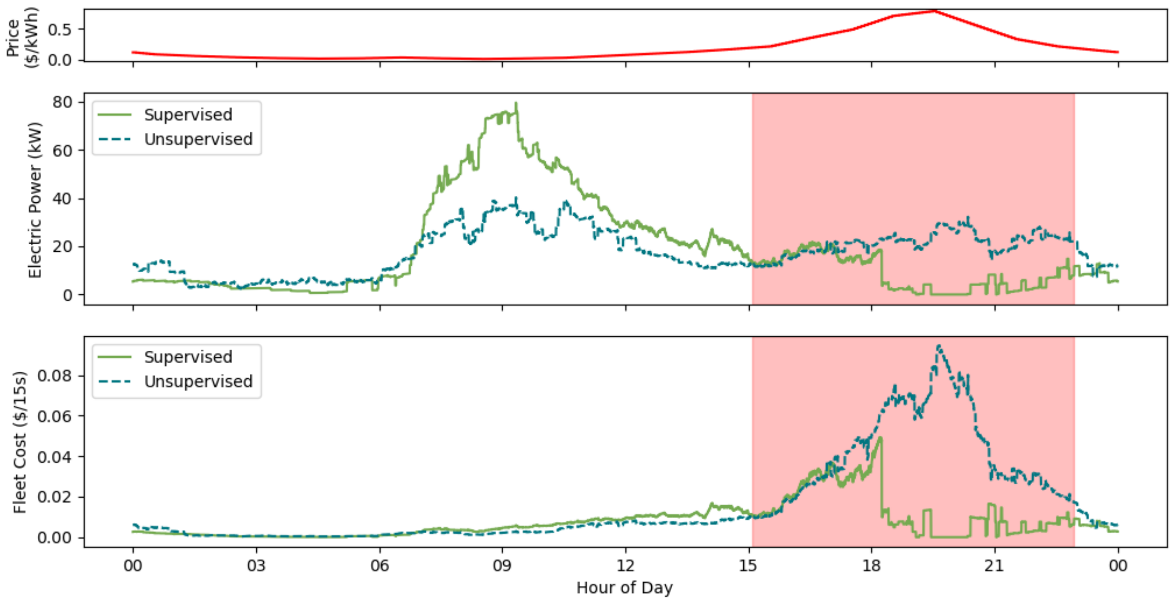

Figure 22. Simulation results with the ALU capability

New Buildings Institute (NBI) is a nonprofit organization driving better energy performance in buildings to make them better for people and the environment. We work collaboratively with industry market players—governments, utilities, energy efficiency advocates, and building professionals—to promote advanced design practices, innovative technologies, public policies, and programs that improve energy efficiency.

To learn about the Advanced Water Heating Initiative visit www.advancedwaterheatinginitiative.org/join-us.