

Aggregation of Residential Water Heaters for Peak Shifting and Frequency Response Services

THOMAS CLARKE (Student Member, IEEE), TYLOR SLAY (Student Member, IEEE),
CONRAD EUSTIS^{ID}, AND ROBERT B. BASS^{ID} (Member, IEEE)

Department of Electrical and Computer Engineering, Portland State University, Portland, OR 97201 USA

CORRESPONDING AUTHOR: R. B. BASS (robert.castro@ladwp.com)

ABSTRACT The increased penetration of renewable energy resources poses challenges for grid stability. The stochastic generation of solar and wind power cannot be controlled to follow load. And, the transition away from synchronous generators is reducing the capacity to arrest and recover from frequency disturbances. Smart electric water heaters provide utilities with an appliance that can be remotely controlled and serve as a form of energy storage. They have very fast response times and make up a large amount of residential energy consumption, making them useful for load peak shifting as well as other ancillary grid services. As smart appliances become increasingly widespread, more and more devices can be brought into the utility control network and aggregated into a flexible resource on a multi-megawatt scale. This paper demonstrates the usefulness of aggregated electric water heaters for providing two ancillary services: peak shifting and frequency response. Because a large number of assets are required, emulators are developed based on observations of real devices. Emulated water heaters are then connected to an energy resource aggregator using an internet-of-things network. The aggregator uses these assets to shift consumption away from peak hours and for detecting upward frequency disturbances.

INDEX TERMS Water heaters, DER, demand response, frequency response, peak shifting, aggregation.

I. INTRODUCTION

THE increasing contribution of renewable energy sources to electricity generation comes paired with new challenges for power grid reliability. Traditionally, the vast majority of electric power has been provided by a relatively small group of large, centrally-controlled generators. These generators can be dispatched to operate above or below their normal levels in response to changes in the balance between supply and demand. Additionally, they provide stability to the grid electrical frequency through their mechanical inertia. However, both of these qualities may be lost in the transition to renewable energy sources. Two of the main renewable sources, wind and solar power, are inherently stochastic in their generation capacity. The sun cannot be made to shine, and the wind cannot be made to blow. The times when supply is available often do not match the times of peak consumption, and the climatic factors that influence power output often also affect the load profile, leading to greater unpredictability. Generation sources that provide energy to

the grid through DC/AC or AC/AC power conversion, such as photovoltaics, battery-inverter systems, and Type 4 wind turbine generators, have no mechanical inertia whatsoever, and detract from the grid's resilience to frequency disturbances. Due to the distributed nature of sources such as rooftop solar installations, accidental or intentional disconnection of a load region from the grid results in loss of generation as well, and may exacerbate problems.

Regardless of the transition to renewable generation, some characteristics of the power system will not change. The amount of energy supplied and consumed by the grid must always match. This balance can be maintained through control of generation, control of load, or both. The large-scale storage of readily dispatchable electric power, such as through batteries or pumped-storage hydroelectricity, is not currently economically feasible and possibly never will be [1]. If the new means of generation cannot be controlled sufficiently to provide grid stability, then the solution must lie with real-time control of power consumption. Providing

the grid with energy balancing and ancillary services through control of the load is known as “demand response” (DR). This research characterizes smart electric water heaters as an asset for DR, and tests their use for load peak shifting and recovery from frequency disturbances.

Water heaters are primary candidates for DR, and a large body of work covering their potential application already exists [2]–[4]. The research presented here distinguishes itself by incorporating the individual devices into an aggregator called Distributed Energy Resource Aggregation System (DERAS). DERAS is a custom-built aggregation system that uses an internet-of-things (IoT) framework to concurrently communicate with multiple assets. Each individual asset is equipped with a distributed control system (DCS) that observes and reports parameters such as the asset’s current state and the amount of energy that it can absorb. DERAS aggregates these parameters for all of the assets in its network and provides the user with data describing the entire system. Some benefits of this approach are that numerous small resources can be observed and controlled on a very large scale while preserving customer anonymity. Additionally, the stochastic differences between devices are evened out as more assets are added to the network, yielding more predictable load and generation patterns. By sending control signals to the assets, DERAS can then provide several services, such as executing a predetermined energy-balancing schedule or responding to frequency disturbances.

One of the enabling technologies for this is the CTA-2045 communication protocol specification, which allows DERAS to communicate with a variety of devices from different manufacturers. Two water heater models are used for this study: one with resistive heating elements, and one hybrid with both a heat pump and resistive elements. Each device comes from the manufacturer equipped with a CTA-2045 interface. The objective of this work is to demonstrate that DERAS can provide valuable ancillary services to a utility using a group of water heater assets via CTA-2045, specifically peak shifting and frequency response.

For the remainder of this work, the term ‘EWH’ will be used to refer to traditional resistive water heaters, and ‘HPWH’ will be used for heat pump water heaters.

II. LITERATURE REVIEW

A. THE CTA-2045 COMMUNICATIONS SPECIFICATION

The CTA-2045 specification was developed from 2008 through 2012 by the Electric Power Research Institute (EPRI) and the Smart Grid Interoperability Panel (SGIP), with the goal of creating a universal demand response communications standard for a wide variety of devices and manufacturers [5]. The standard also defines a socket interface so that devices come ready for energy management functions directly from the manufacturer. Utilities can design and sell universal communication modules (UCMs) that plug into the CTA-2045 interface. The utility is then able to exchange telemetry with the UCM using whatever communication

system they desire, such as Wi-Fi, radio, cellular networks, or even systems that haven’t yet been invented [6]. The UCM translates that telemetry to and from the CTA-2045 standard. This flexibility is critical for appliances with long service lives, because major developments and changes to communication networks can take place over the device’s lifespan. The cost of equipping devices with UCMs is expected to decline steeply as CTA-2045 becomes more widespread. A study by the Bonneville Power Administration (BPA) has projected the cost of equipping a single water heater with a UCM dropping to \$25 by 2030 [5]. Beyond the benefit of allowing utilities to design whatever UCM is most practical for them, the modular approach also has cybersecurity advantages. If vulnerabilities in the interface are exposed, the UCM is much cheaper and easier to update or replace than the appliance’s built-in hardware.

B. WATER HEATERS FOR DEMAND RESPONSE

Utilities have been making use of DR in the industrial sector for years by engaging the largest energy consumers through dynamic tariffs or incentivized direct load control [1], [7]. However, the residential sector’s DR resources remain largely untapped. The residential sector constitutes 37% of electricity consumption in the U.S. [8] Water heating makes up a significant portion of this - 17% of the national residential consumption [9], and up to 30% in some regions [10]. Additionally, the times of water heater usage coincide with the peaks in overall electricity demand, meaning that control of water heaters can be a powerful tool for peak shifting.

Water heaters are valuable DR assets for several more reasons: First, the rated power consumption of an individual device (typically 4500 W) is high, meaning that relatively few devices can quickly accumulate to have a bulk-scale impact. Second, EWHs turn on and off very quickly. The ramp rate from zero to full rated power is effectively instantaneous, allowing for more precise control and response to short-term transient events such as frequency disturbances. Third, resistive water heaters are almost purely resistive and do not require reactive power support from the grid (unlike HVAC systems or induction machines, for example). Finally, water heaters serve as a form of energy storage, and the time of hot water use is not highly correlated to the time of electricity consumption.

The usefulness of water heaters for DR increases with the number of available assets, and is therefore directly dependent on customer participation and the market penetration of accessible devices. A pilot study of CTA-2045-equipped water heaters conducted by the BPA projects that, if 26.5% of all electric water heaters in the states of Oregon and Washington were enrolled in a DR program, the combined dispatchable load would be equivalent to a 301 MW peaking plant. This resource has an estimated long-term value of \$106 million, and on a national scale extrapolates to \$2 billion [5]. The report claims a benefit-cost ratio of 1.0 even if customer participation is as low as 5%. The BPA also proposes a market plan for utilities and the three major electric

water heater manufacturers (AO Smith, Rheem, and Bradford White) to fill the Oregon and Washington markets with CTA-2045-equipped water heaters starting in 2025. Given the expected life-spans of traditional water heaters already in use, the report expects 91% of all water heaters in the region to have CTA-2045 interfaces by 2039.

Another challenge to residential DR is recruiting and maintaining a large number of participating customers. The primary concern for water heaters is that shed events will cause customers to run out of hot water, but this was only a minor issue in the BPA pilot study. The BPA enrolled 277 households, and sent approximately 600 DR commands over the course of 220 days. Customers received different compensation based on their local utility. Portland General Electric (PGE) offered customers a total of \$50 for starting and \$100 for completing the pilot [5]. Customer feedback was majority positive; 80% of customers reported that they were very satisfied with the program, and 94% stated that they would be very likely to enroll in similar programs in future. Only two households left the program due to hot water shortages.

Although HPWHs are more energy-efficient than resistive EWHs, they are less useful from a DR point of view. They aren't capable of absorbing as much energy from the grid, they take much longer to turn on, and the rated power of a heat pump is usually much less than the rated power of a resistive heating element. The BPA points out that, even though EWHs will only make up 69% of water heaters in 2019, they will still represent 80% of the utility benefits. Additionally, because HPWHs reheat more slowly, utilities must be more conservative when shedding them to avoid hot water shortages [5].

C. FREQUENCY RESPONSE SERVICES

The ongoing retirements of synchronous generators around the country and the increase in power provided by sources with little or no mechanical inertia erode system frequency response. While NERC's 2018 long-term reliability assessment predicts that the national grid frequency response resources are expected to remain adequate through 2022, the large-scale transitions to renewable resources may put the system in serious jeopardy in the following decades. Type 1, 2, and 3 wind turbines use induction generators and are not mechanically synchronous with the grid electrical frequency. As such, they cannot provide governor response in the same way as more traditional power plants. Type 4 wind turbines, which connect to the grid through electronic AC/AC power conversion, and DC energy sources that require inverters, such as battery systems and PV, do not naturally provide any inertia to the grid.

Frequency response capacity has a high monetary value. The Salem Smart Power Center (SSPC) is a PGE R&D project located in Salem, Oregon, and is home to a 5 MW, 1.25 MWh battery energy storage system (BESS). One of the goals of the SSPC is to determine the value of different ancillary services, including frequency response. While the SSPC only engaged in frequency response for a total of 17 hours in

a full year, it was still found to be more valuable than all other services combined [11].

While NERC requires utilities to provide frequency response services, they do not specify when an event is occurring, or exactly how the response must be executed. The SSPC's event detection and response algorithms were created by PGE, and only react to negative frequency deviations. As soon as an event is detected, the battery output ramps up as quickly as possible to nearly full power. The power remains at this level for 3 minutes, before slowly ramping down as the Secondary and Tertiary frequency control resources take over. 300 kWh (24%) of the SSPC battery's storage capacity is reserved for frequency response at all times, which is roughly equal to the total energy discharged during a response.

Battery inverter systems are one of the best options for providing frequency response to negative frequency deviations because they can maintain an energy reserve to discharge to the grid whenever needed. Water heaters cannot generate power, but they can change the overall load. In the event of a negative frequency deviation, water heaters would need to turn off in order to provide frequency response. However, because water heaters are already off for most of the time [12]–[14], a very large number of devices would be required to provide a significant response to a frequency decline. However, EWHs are an ideal resource for responding to upward frequency disturbances. Caused by a sudden loss of load or spike in generation, these disturbances are less common but pose similar threats to grid reliability as downward disturbances. Arresting the positive frequency change requires a decrease in generation or increase in load. Because they ramp up to their full power in sub-cycle time frames, EWHs are capable of responding to these events even faster than battery inverter systems. Additionally, if this service is delegated to EWHs, battery inverter systems on the same network won't need to maintain any charging headroom.

Frequency disturbance events are fairly recognizable after they have occurred. However, effective frequency response requires that events be detected within the first few seconds. This requires service providers to create an automated event detection algorithm. Because the grid's generation/load balance is in a constant state of change, a detection algorithm that is too sensitive may generate many false positives. For example, the SSPC successfully responded to 15 out of 18 registered frequency disturbances over the course of 10 months in 2016, equating to an 83.3% detection rate, but also triggered erroneously nearly eight times as often [11]. Responses to these false positives are both a waste of resources and an additional stress on the system. A balance must be struck in order to minimize both false negatives and false positives. NERC provides little guidance for detecting frequency events in real time, but does establish a minimum stable frequency for each Interconnection. If the grid frequency does not drop below this "floor frequency," there is no need for a frequency response. In the Western Interconnection, this frequency is 59.976 Hz. Frequency response measures should be able to

detect and arrest a frequency decrease somewhere between 59.976 Hz and 59.5 Hz, where the last-resort measure of under-frequency load shedding (UFLS) begins [15]. NERC generally refers to frequency changes as \pm deviations from 60 Hz, so equivalent thresholds can be applied to upward frequency disturbances.

III. METHODS

This work seeks to evaluate the usefulness of an aggregation of water heaters in providing demand response and frequency response services. By nature, a large number of devices is required [14]. Conducting a study such as the BPA pilot project requires the participation of utilities and a large number of their customers, which far exceeds the means and scope of this research. Additionally, the internet-of-things network that is used to exchange telemetry between the aggregator and devices is a novel approach that has not yet expanded beyond a local area network (LAN). Obtaining, installing, and operating a large number of water heaters would also be very expensive and labor intensive, require an extravagant amount of space, and consume a large amount of water.

For these reasons, virtual devices are used primarily here, rather than physical devices. Developing these virtual device models does require the characterization of physical devices, and a test station was built for that purpose. A thermal model for a small electric water heater was developed by a previous study [16]. Part of the work presented here builds on that original model and expands it to larger resistive and heat-pump water heaters.

Evaluation of the CTA-2045 interface is also necessary, both as a means of controlling and monitoring the devices, and because the exchange of telemetry is limited by the interface capabilities and functionality. Once emulators have been developed that adequately reflect the characteristics and behavior of the physical devices, a large number can be run and connected to the aggregator simultaneously. The aggregator is unaware that these devices are virtual, and they send the same information that a CTA-2045 interface would provide.

The starting energy condition and usage schedule of each virtual device is randomized to simulate the stochastic conditions of real appliances. The net power consumption and energy capacity of the virtual devices can then be used to evaluate the effects of demand response and frequency response efforts by the aggregator. In this study, the goal for demand response is to shift power consumption away from when it would normally occur [17]. The goal for frequency response is to detect upward frequency disturbances and react accordingly by changing the net load.

A. DEVICE RESPONSES TO CTA-2045 COMMANDS

One of the purposes of the DCS is to convert instructions from the aggregator into the most appropriate CTA-2045 command. As far as the aggregator is concerned, water heaters are assets that can be either importing their full rated power (on) or not (off). Relevant instructions for water heaters therefore

TABLE 1. Available import energy thresholds (watt-hours) under different CTA-2045 commands.

	EWH		HPWH	
	Low	High	Low	High
Baseline	0	900	0	1200
Shed	1725	2100	75	1575
Critical Peak Event	2325	2475	675	1800
Load Up	0	300	0	150

come as a request to either turn on or turn off. The asset's ability to follow these instructions is dependent on its current energy state and the temperature thresholds set by the manufacturer. Every water heater has a default regulation range bounded by upper and lower thresholds. The stored energy bounces back and forth between the thresholds as the tank is heated and water is drawn.

CTA-2045 commands effectively switch the water heater between different sets of thresholds, therefore shifting and/or changing the size of the regulation range. Because the water heater will always try to keep its stored energy inside of the regulation range, adjusting the regulation range can result in the device turning on or off. Ideally, it would be possible to set very narrow regulation ranges at either extreme of the allowable energy range, which would allow the most precise control of the device. The total range over which it is possible to turn the device on and off freely is referred to as the 'dispatchable' range from this point on. This dispatchable range is limited by the absolute limits on acceptable water temperature set by the manufacturer, and by the thresholds of the different regulation ranges available through CTA-2045 commands. Keeping the dispatchable range as wide as possible maximizes the aggregator's utility by increasing the number of devices that it is able to control at any given time. Additionally, determining the thresholds for the different CTA-2045 commands is necessary in order to properly model device behavior.

The thresholds for each CTA-2045 command can be observed by logging the available import energy (how much more energy the device can import before being fully heated, from here on referred to simply as 'import energy') over a period of time while drawing hot water occasionally. These thresholds were observed to be the same regardless of the temperature setpoint, so all tests were conducted with a setpoint of 120 °F. To avoid possible correlations between usage pattern and water heater behavior, randomized draw schedules were used. The observed thresholds for the CTA-2045 commands *shed*, *load up*, and *critical peak event*, as well as the baseline thresholds when no commands are issued, are summarized in Table 1.

Most of the regulation ranges are much wider for the HPWH than for the EWH, meaning that control of the HPWH is more limited. Additionally, the *shed* command does not effectively turn off the heat pump of the HPWH unless it is nearly fully heated. The *critical peak event* does a better job of controlling the HPWH, but this command is intended to be used infrequently.

TABLE 2. Peak shifting schedule used for emulator and demand response validations.

Time Window	Command
Midnight - 6AM	Load up
6AM - 11AM	Shed
11AM - 4PM	Load up
4PM - Midnight	Shed

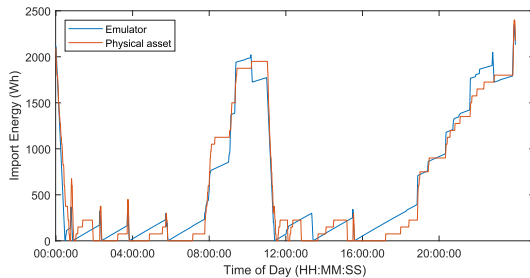


FIGURE 1. EWH emulator validation. Behavior of the emulated EWH (blue) closely matched that of the physical device (orange).

B. EMULATOR VALIDATION

Programming emulators that faithfully represent the physical devices required modeling of five main observable dynamics:

- 1) Convection losses
- 2) Change in stored energy from usage events
- 3) Power consumption and heating
- 4) Device behavior (when the resistive elements and/or heat pump turn on and off)
- 5) Response time to CTA-2045 commands

For validation, the emulator and physical device were run with the same schedule of usage events and CTA-2045 commands. The usage schedule was obtained using a probabilistic tool for generating realistic water heater usage schedules created by NREL [12]. The CTA-2045 commands follow a generic peak-shifting dispatch schedule summarized in Table 2.

A side-by-side comparison of the emulator and physical EWH is shown in Figure 1. The emulator is a close, but not a perfect match for the physical device. The goodness of fit, determined using a normalized root mean square error cost function, is 78.66%. While this match is less than ideal, particularly in the period between noon and 4PM, the physical device is also not particularly consistent from day to day. The difference between the physical and emulated devices' energy states is never greater than a few hundred watt-hours, so the emulators were considered to be good enough to substitute for physical devices.

C. DEVICE RANDOMIZATION

The emulator programs used for aggregate testing take several measures to provide the stochastic differences in usage expected from real customers. The first measure is accounting for household size. The NREL water heater profile generation tool was used to create separate schedules for houses with one

to five bedrooms. The U.S. Census reports that, as of 2017, 13% of households had no bedrooms or one bedroom, 26.3% had two bedrooms, 39.6% had three bedrooms, 16.5% had four bedrooms, and 4.4% had five or more bedrooms [18]. Each time an emulated device is created, it receives a schedule based on the probabilities listed here using a random number generator. The second measure is randomization of the starting condition. The initial import energy of a new device is determined using a normal distribution with a mean of 50% and standard deviation of 30% of the rated value. The third measure is randomization of the usage schedule. Both the times and volumes of use are spread over normal distributions. For each usage event, the mean for the time distribution is the base value, and the standard deviation is one half-hour. For the volumes, the mean is the base value, and the standard deviation is 30% of that.

D. AGGREGATOR

DERAS, the aggregator used in this work, was first developed and presented in [19]. DERAS is a C++ program that uses Alljoyn, an open source software framework for creating networks of different devices. The user can query information on the total aggregate or on individual devices, directly request the aggregate to begin importing or exporting a specific amount of power, or begin different services. A service for scheduled dispatch commands was previously developed [19]. The frequency response service was developed for this research.

E. DIGITAL TWINS

In order to include a large number of assets on the network, it is necessary to limit the amount of network traffic created by each device. To accomplish this, DERAS uses digital twins to represent devices in its network. Each time a device is added to the network, DERAS creates a digital twin for it locally, which serves as a placeholder for the asset's energy and power properties. For example, if an EWH joins the network, DERAS queries its rated power and energy capacity as well as its current state, and stores all of these values in a new digital twin. If the EWH is instructed to be on, DERAS automatically begins updating the energy value stored in the digital twin using the rated power. Digital twins also include an idle loss rate. The digital twin is not intended to accurately track the state of the asset over a long period of time, but rather only to reduce the necessary frequency of metrology requests. For this research, the DCSs were programmed to send property updates once every hour, or any time the device import energy changed by more than 10% of its rated value. In this way, any sudden changes to the device energy state, such as a large usage event, are soon reported to the digital twin. However, for most of the time, when very little is happening, the device only needs to send an update once every hour. This dramatically reduces the total network traffic necessary for a fairly accurate real-time assessment of the aggregate. A validation test with a physical EWH demonstrated that the digital twin never strayed from the device's reported energy

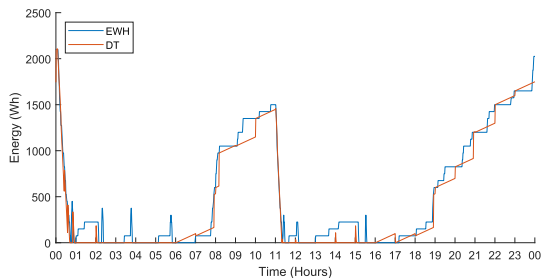


FIGURE 2. Comparison of physical EWH behaviour (blue) to that of its Digital Twin (orange).

values by more than a few hundred watt-hours, as shown in Figure 2.

Using the available machines and software, it was possible to include about 100 assets in the DERAS Alljoyn network. Adding more assets made the system increasingly unstable, resulting in communication timeouts and disconnections. Therefore, all tests for this research were performed using 100 assets.

F. PEAK SHIFTING SERVICE

To test the load-shifting capacity of an aggregate of water heaters, the DERAS peak shifting service was used. This experiment used the same generic 24-hour dispatch schedule used in the emulator validation tests and shown in Table 2, with the goal of shifting power consumption away from the traditional peak hours in the morning and evening. First, a baseline test was run to show the pattern of the aggregate if no commands were sent. The aggregator was then programmed to send the scheduled commands to all devices in its network, staggering the *load up* command and prioritizing colder heaters to avoid sudden spikes in power consumption from all devices coming on at once. Because [5] predicts a 31% market penetration of HPWHs by the time that CTA-2045 interfaces are essentially ubiquitous, the mixed test was done with an aggregate of 69 EWHs and 31 HPWHs.

G. FREQUENCY RESPONSE SERVICE

The DERAS frequency response service consists of two algorithms: one for detecting positive frequency disturbance events, and one for responding to them. In actual use, DERAS would receive second-by-second updates of the grid frequency in order to detect events. For this research, DERAS was provided with a data file containing actual frequency data from previous disturbances provided by PGE.

The detection algorithm continuously monitors frequency and calculates slew rate using a sliding window. A positive frequency event is declared if a slew rate greater the 0.3 mHz/s is detected *and* if the frequency deviates above 60.025 Hz.

This work treats water heaters as upward dispatchable resources, which are normally idle but can be turned on. Once an event is detected, DERAS initiated the response algorithm. Because frequency response is a high-value and time-critical

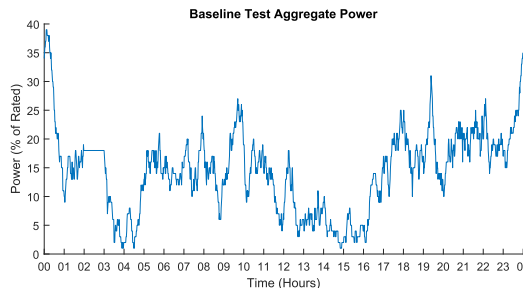


FIGURE 3. Power Consumption of Aggregate without any DERAS Commands.

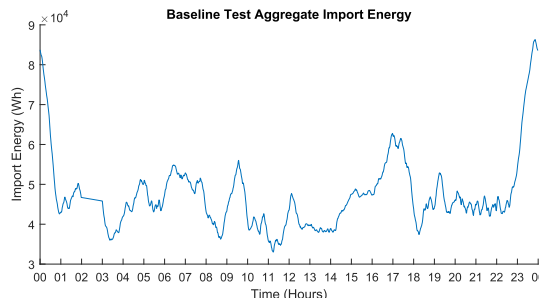


FIGURE 4. Import Energy of Aggregate without any DERAS Commands.

service, the aggregator initially responds to an events with its full capacity. All possible devices will be turned on as soon as possible. After three minutes of all devices in the aggregate importing as much power as they can, the aggregator ramps down the power to zero over the course of another three minutes.

IV. RESULTS

A. PEAK SHIFTING TEST

To demonstrate the utility of DERAS for peak shifting using EWHs, an aggregate of 100 EWH emulators was first run with no CTA-2045 commands to show the power and energy trends with no outside influence. The baseline power consumption is shown in Figure 3.

The typical residential water heater load curve is discernable within Figure 3, with high consumption in the morning and evening, and low consumption in the afternoon and at night [20], [21]. Though, the power profile is noisy due to the small number of assets used in this study. The greatest consumption is around midnight, when nearly 40% of the aggregate’s devices are on. This spike at midnight is anomalous and most likely due to the randomization of the usage schedules. The import energy for the same aggregate is shown in Figure 4.

Aside from a peak during the biggest usage hour leading up to midnight, the baseline import energy does not follow any particular pattern, but generally stays within a band between 40 and 50 kWh. This meets expectations, because the devices are all staying within the same regulation range for the entire test. The baseline regulation range is from 0 to 900 Wh, which

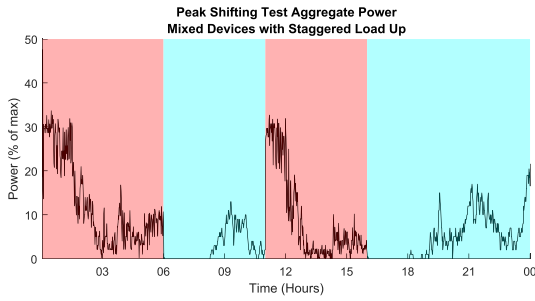


FIGURE 5. Aggregate Power of Mixed Devices with Staggered Peak-Shifting Schedule.

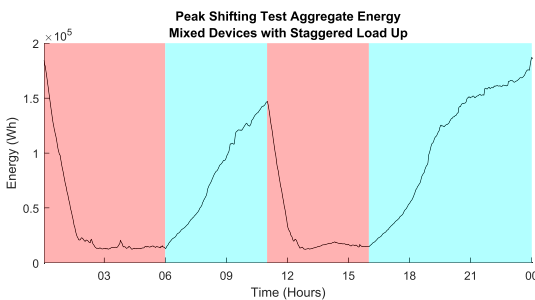


FIGURE 6. Aggregate Import Energy of Mixed Devices with Staggered Peak-Shifting Schedule.

for 100 devices aggregates to a range of 0 to 90,000 Wh. If the aggregate import power averages to around 45,000 Wh, this puts it right in the middle of the regulation range.

In the next test, DERAS controlled 69 emulated EWHs and 31 emulated HPWHs over 24 hours using the generic dispatch schedule described in Table 2 with its peak shifting service. The resulting aggregate power and energy plots are shown in Figures 5 and 6. Load up times are shaded red, and shed times are shaded blue.

There is a clear reduction in power consumption during the shed hours. Particularly at the start of the peak usage hours, from 6 AM to 9 AM and from 4 PM to 9 PM, there is a significant reduction in power consumption because the devices begin with a large amount of energy and widen their regulation ranges. The effectiveness of the evening shed period does gradually expire as it gets closer to midnight, as the aggregate's amount of stored energy becomes exhausted and more devices turn back on. The maximum effective length of a shed period therefore appears to be around 4 hours, although this will always depend on usage patterns.

B. FREQUENCY RESPONSE TEST

To test DERAS's capacity to detect and respond to positive frequency disturbances, frequency data containing an event were applied to an aggregate of 100 emulated EWHs. The response to this event is shown in Figure 7.

As shown in Figure 7, the event is detected during the ramp period of the event, and the response occurs very quickly. Only four of the devices were on before the event, but all

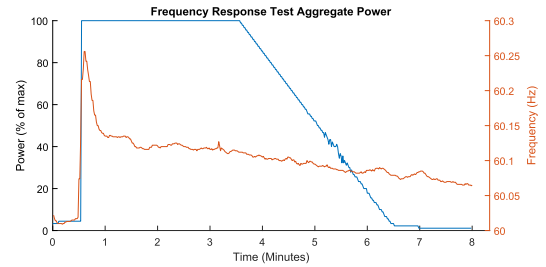


FIGURE 7. Aggregate Frequency Response to Upward Disturbance.

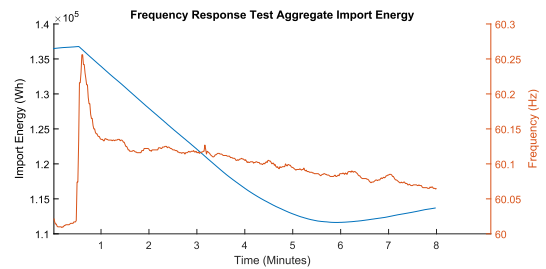


FIGURE 8. Aggregate Import Energy during Frequency Response.

100 were on within seconds of the event beginning. Given the fast *load up* response time for EWHs, this test indicates that a very large number of EWHs could be dispatched in the first few seconds of a frequency disturbance, even considering the delay between the DERAS command and the physical response. All 100 devices are able to remain on for the desired 3 minutes, and then ramp down back to pre-disturbance conditions at the end of the following three minutes. The response demonstrated here is a successful recreation of the SSPC's frequency response, modified for upward frequency disturbances. This work does not suggest that water heaters alone are sufficient for frequency response. The Western Interconnection is a massive system and no single asset is capable of such an impact on the grid frequency. This work only proposes that aggregated water heaters can contribute to the system's net resilience to upward frequency disturbances. The effect of this response on the aggregate's available import energy is shown in Figure 8.

Over the course of the response, the aggregate's import energy was decreased roughly from 136 kWh to 112 kWh, or by about 24 kWh. An aggregate of 100 EWHs will almost always have this much available import energy unless it is already loading up. Because the power consumption is also proportional to the number of devices in the aggregate, it is safe to say that an EWH aggregate of any size can be expected to respond to upward frequency disturbances with its full capacity most of the time.

The frequency response experiment was repeated with a mixture of 31 HPWHs and 69 EWHs to demonstrate that HPWHs are a less useful asset than EWHs for frequency

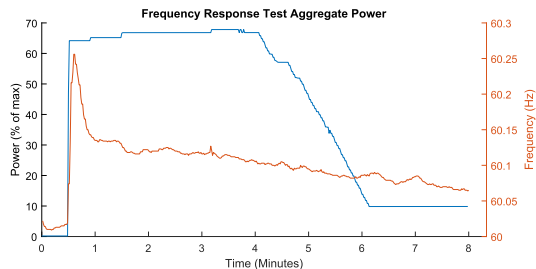


FIGURE 9. Aggregate Power during Frequency Response with Mixed Devices.

response due to their lower power consumption and longer response times. The results of this test are shown in Figure 9.

Replacing some of the EWHs with HPWHs has clearly visible effects on the frequency response. The heat pump emulator was programmed with a one-minute turn-on delay based on observations of the physical device. Therefore, the initial spike in power at the beginning of the disturbance consists only of the EWHs and possibly the resistive elements of some HPWHs which had very high import energy (HPWHs only respond to *load up* commands with their resistive elements if they have more than 1425 Wh of import energy). Two heat pumps are on before the event begins, and 64 elements are energized in the initial response. This means that not all of the 69 EWHs are able to turn on, which is most likely due to them already being too warm to respond to the *load up* command. The heat pumps turn on one minute later, and slowly ramp up their power consumption over the course of the response. However, because the heat pumps consume less than 10% of a resistive heating element's power, even when all 31 heat pumps are on this only makes a small difference to the total response power. The ramp down is also less smooth because the different devices are consuming different amounts of energy. Some usage events occurred during the response, so at the end nine devices are unable to turn off (this is anomalous and could also have happened in the test with only EWHs). In general, replacing EWHs with HPWHs reduces the total power available for frequency response as well as the precision of DERAS's control.

V. DISCUSSION

A. SUMMARY OF RESULTS

The peak shifting and frequency response capabilities of aggregated water heaters demonstrated in this research met all of the initial goals and expectations. Satisfactory emulators were created for CTA-2045-equipped EWHs and HPWHs, and groups of 100 emulated devices were controlled using the custom-built DERAS aggregation system. The tests were conducted both exclusively with EWHs and with a mixture of EWHs and HPWHs.

The peak shifting tests demonstrated that it is possible to displace significant power consumption during periods of high usage, although the duration of this displacement is typically limited to 4 or 5 hours. In initial experiments, the control signal to turn on at the end of a displacement period was sent to all devices simultaneously, resulting in

a sudden, severe spike in consumption. However, this was corrected by adding a limit to the percentage of devices in the network that received the command at once. DERAS prioritizes devices with more available energy, so the coldest heaters come on first.

The frequency response tests demonstrated both that the aggregator's frequency disturbance detection algorithm is capable of quickly triggering a response to disturbances, and that EWHs have great potential to respond to upward frequency disturbances due to their very fast activation times, ramp rates, and high power consumption. An aggregate of EWHs was able to provide a frequency response nearly identical to examples from the SSPC. HPWHs were shown to be less useful for frequency response because of their low power consumption and the necessary turn-on delay of heat pumps. Frequency response is a high-value ancillary service, and dedicating water heater assets to responding to upward frequency disturbances could be extremely beneficial to utilities while minimally impacting customers.

B. POTENTIAL SOURCES OF ERROR

The experimental design had several flaws. One shortcoming is the small number of devices used in the aggregate tests. Aggregate sizes were limited to 100 devices because the Alljoyn network would otherwise become unstable and begin losing devices. Alljoyn has already been absorbed by a larger project and lost most of its support, so the only solution for this problem is to switch to a new, more robust framework.

Another weakness in the experimental design is the small group of profiles used to simulate stochastic customer usage. Only one original profile was made for each household size, and then the emulated devices created unique profiles by randomizing the original ones. This led to a lack of overall variety in the aggregate usage patterns as well as some anomalous results such as the large peak at midnight seen clearly in Figure 3. Diversifying the group of original usage profiles would be an improvement to the experimental design.

VI. CONCLUSION

This work successfully demonstrated the utility of DERAS for peak shifting and frequency response using aggregates of water heaters. These services have high monetary value to the current grid, and will increase in value as renewable energy sources contribute more of the total generation. Turning large numbers of water heaters off at peak hours can help even out the balance of supply and demand and reduce the ramping requirements of the worsening duck curve on the rest of the system. Turning water heaters on in response to upward frequency disturbances can help replace the mechanical inertia of traditional generators as these machines are replaced with inverter-based generation.

This research leaves several opportunities for future work. One possible development is the use of machine learning for frequency disturbance detection. As demonstrated by both the SSPC detection algorithm and the algorithm developed for DERAS, accurately detecting frequency disturbances without

responding to false positives is very difficult. Because a large corpus of frequency data with many identified disturbances already exists, this is an excellent application for guided machine learning.

Increasing the number of devices that can be included in the aggregator's network would improve the reliability of experiments and is also a necessary step for expanding this work into a real-world application with thousands of customers. Therefore, an immediate and pressing task is the transfer of the aggregator and device controllers to a new communication framework. The creators of DERAS are currently building new versions of DERAS and the DCS that will use the IEEE 2030.5 smart energy profile application protocol.

REFERENCES

- [1] E. Koliou, C. Eid, J. P. Chaves-Ávila, and R. A. Hakvoort, "Demand response in liberalized electricity markets: Analysis of aggregated load participation in the German balancing mechanism," *Energy*, vol. 71, pp. 245–254, Jul. 2014.
- [2] D. Lew *et al.*, "The power of small: The effects of distributed energy resources on system reliability," *IEEE Power Energy Mag.*, vol. 15, no. 6, pp. 50–60, Nov./Dec. 2017.
- [3] J. Kondoh, N. Lu, and D. J. Hammerstrom, "An evaluation of the water heater load potential for providing regulation service," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, Jul. 2011, pp. 1–8.
- [4] S. A. Pourmousavi, S. N. Patrick, and M. H. Nehrir, "Real-time demand response through aggregate electric water heaters for load shifting and balancing wind generation," *IEEE Trans. Smart Grid*, vol. 5, no. 2, pp. 769–778, Mar. 2014.
- [5] *CTA-2045 Water Heater Demonstration Report*, Bonneville Power Admin., Portland, OR, USA, Nov. 2018.
- [6] *Introduction to the ANSI/CEA-2045 Standard*, Electr. Power Res. Inst., Palo Alto, CA, USA, May 2014.
- [7] *Assessment of Industrial Load for Demand Response Across U.S. Regions of the Western Interconnect*, Oak Ridge Nat. Lab., Oak Ridge, TN, USA, Sep. 2013.
- [8] *Electric Power Annual 2017*, U.S. Energy Inf. Admin., Washington, DC, USA, Dec. 2018.
- [9] J. J. Guo, J. Y. Wu, and R. Z. Wang, "A new approach to energy consumption prediction of domestic heat pump water heater based on grey system theory," *Energy Buildings*, vol. 43, no. 6, pp. 1273–1279, 2011.
- [10] R. Diao, S. Lu, M. Elizondo, E. Mayhorn, Y. Zhang, and N. Samaan, "Electric water heater modeling and control strategies for demand response," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, Jul. 2012, pp. 1–8.
- [11] *The Salem Smart Power Center: An Assessment of Battery Performance and Economic Potential*, Pacific Northwest Nat. Lab., Washington, DC, USA, Sep. 2017.
- [12] B. Hendron, J. Burch, and G. Barker, "Tool for generating realistic residential hot water event schedules," Nat. Renew. Energy Lab., Golden, CO, USA, Tech. Rep., Aug. 2010.
- [13] M. Pipattanasomporn, M. Kuzlu, S. Rahman, and Y. Teklu, "Load profiles of selected major household appliances and their demand response opportunities," *IEEE Trans. Smart Grids*, vol. 5, no. 2, pp. 742–750, Mar. 2014.
- [14] K. Marnell, "Large-scale DER aggregations of electric water heaters and battery inverter systems," M.S. thesis, Portland State Univ., Portland, OR, USA, Jun. 2019.
- [15] *Frequency Response Standard Background Document*, North Amer. Electr. Rel. Corp., Atlanta, GA, USA, Nov. 2012.
- [16] A. Clarke, "Electric water heater modeling for CTA-2045 demand response testing," M.S. thesis, Portland State Univ., Portland, OR, USA, Jun. 2018.
- [17] T. L. Clarke, "Aggregation of electric water heaters for peak shifting and frequency response services," M.S. thesis, Portland State Univ., Portland, OR, USA, Jun. 2019.
- [18] *2013-2017 American Community Survey 5-Year Estimates: DP04—Selected Housing Characteristics*, United States Census Bur., Suitland-Silver Hill, MD, USA, 2018.
- [19] T. Slay, "Adoption of an Internet of Things framework for distributed energy resource coordination and control," M.S. thesis, Portland State Univ., Portland, OR, USA, Jun. 2018.
- [20] M. H. Nehrir, R. Jia, D. A. Pierre, and D. J. Hammerstrom, "Power management of aggregate electric water heater loads by voltage control," in *Proc. IEEE Power Eng. Soc. Gen. Meeting*, Jun. 2007, pp. 1–6.
- [21] *Description of Electric Energy Use in Single Family Residences in the Pacific Northwest, 1986–1992*, Office Energy Resour., Bonneville Power Admin., Portland, OR, USA, Dec. 1992.