

# Resource Study of Large-Scale Electric Water Heater Aggregation

KEVIN MARNELL<sup>1</sup> (Member, IEEE), CONRAD EUSTIS<sup>1,2</sup>, AND  
ROBERT B. BASS<sup>1,3</sup> (Member, IEEE)

<sup>1</sup>Pacific Power, Portland, OR 97232 USA

<sup>2</sup>Portland General Electric, Portland, OR 97204 USA

<sup>3</sup>Maseeh College of Engineering and Computer Science, Portland State University, Portland, OR 97201 USA

CORRESPONDING AUTHOR: R. B. BASS (Robert.bass@pdx.edu)

This work was supported by Portland General Electric.

**ABSTRACT** Residential-scale distributed energy assets, like residential electric water heaters, individually present a negligible load to the power grid. When aggregated, however, these assets can impart significant effects within a balancing area; they may be dispatched *en masse* to provide grid services. An aggregation of water heaters may be controlled to assume generator-like functions with the ability to effectively “decrement power” through dispatch of load. This resource study examines the capabilities of a 10,000 unit water heater aggregation by subjecting the aggregate to dispatch requests of various size and duration, then analyzing how the aggregate responds to and recovers from these requests. Results show that a large-scale aggregation of electric water heaters may effectively decrement power on the scale of megawatts when the dispatch request size and duration are appropriately considered.

**INDEX TERMS** Demand response, distributed energy resources, DERMS, aggregation, CTA-2045, electric water heaters, energy take.

## I. INTRODUCTION

**D**ISTRIBUTED Energy Resources Management Systems (DERMS) have been recognized as a means to provide energy or manage loads to better operate the power grid [1]. By injecting power into the grid or reducing load, DERMS can accomplish a number of useful grid functions. In the most basic of terms, when generation is greater than load, DERMS can turn on additional loads to reestablish the balance. Likewise, when load is greater than generation, DERMS can turn off loads. These simple dispatch commands may be used to provision specific ancillary grid services such as frequency regulation, spinning and non-spinning reserve, and peak demand mitigation [2].

Aggregations of Electric Water Heaters (EWH) could be managed by a DERMS to fulfill energy bids. Laurent and Malhame developed a model to consider the physical characteristics of EWHs including insulation, element rating, and water demand [3]. They used this model to test the aggregate behavior of the EWHs after power interruptions. Fitzgerald *et al.* simulated an aggregation of 100,000 EWHs to improve efficiency of wind generation, observing a decrease of 25% in electric power usage and a 38% decrease

in energy costs [4]. Roux *et al.* developed a peak demand manager algorithm that encouraged EWHs to compete for use of the grid during peak hours, limiting electricity to EWHs with the greatest need [5]. Li *et al.* simulated an aggregator with EWHs, solar panels, battery energy storage, and electric vehicles intended for large multi-tenant buildings [6]. They developed control algorithms for the resources to minimize power costs through load shifting and energy trading, leading to shorter return on building investment and lower energy bills for tenants. Kapsalis and Loukas simulated a computerized optimal scheduling algorithm that aimed to keep electrical costs low while minimizing customer discomfort [7]. They assumed control of the setpoint features of their EWH aggregation. Kapsalis *et al.* presented a heuristic scheduling algorithm to balance decisions between consumer comfort and power costs [8]. They simulated their algorithm in a real-time energy market and found it performed as well as a standard computerized optimization model with considerably less overhead.

EWHs have been harnessed by utilities via large-scale DERMS programs. Utility load control projects involving residential EWHs were implemented as far back as the late

1970s when Florida Power and Light installed controllable EWHs in customer homes in Boca Raton [9]. By the 1990s, this program had grown to include hundreds of thousands of customers. More recently, Great River Energy, an electrical cooperative in rural Minnesota, implemented a residential EWH aggregation system to provide demand response [10]. Currently, Great River Energy manages an aggregation of over 110,000 EWHs, amounting to almost 20% of its customer base. Hledik *et al.* examined the use of on/off switches to control the Great River Energy aggregation, which includes a mixture of resistive water heaters and heat pump water heaters [11]. The Pacific Northwest GridWise demonstration project conducted a field demonstration in communities of the Olympic Peninsula [12]. Participants provided access to a variety of residential loads, including water heaters. These were used to manage feeder congestion through peak load reduction using price signals via a two-way communication system.

The purpose of this manuscript is to investigate the resource availability of a large aggregation of EWHs, to subject these aggregated EWHs to dispatch requests of various bid sizes and duration, and to analyze how the aggregation responds to, then subsequently recovers from, a dispatch. The EWH aggregation model uses dispatch commands that conform to the industry-developed ANSI/CTA-2045 communication standard, developed by the Consumer Technology Association [13].

The remainder of this paper is organized as follows. Section II discusses the methods and methodology used to build this EWH aggregation and simulate its behavior to energy bids. Section III presents the results of the aggregator simulation to various energy bids. Section IV discusses the results of the simulation. In Section V, the authors provide concluding remarks [14].

## II. EXPERIMENT METHODS

### A. AGGREGATION MAKEUP

A model of 10,000 EWHs with unique hot water draw profiles was developed. The water draw profiles were produced using the Domestic Hot Water (DWH) Schedule Generator created by Hendron *et al.* for the U.S. Department of Energy [15]. This event generator randomly produces a schedule of hot water draws for households in the United States. The event list is built using macros within Excel. The key inputs to the event generator are setpoint of the water heater, number of bedrooms, and a U.S. city from which weather data are gathered.

The behavior of water heaters depends on household size. U.S. census data were used to apportion a population of one, two, three, four, and five bedroom households [16]. Households noted to have “zero bedrooms” in the Census data were assumed to behave like households with one bedroom. Where the Census referenced “five or more” households, these were assumed to all be five bedroom households.

For this simulation, all water heaters were given a standard setpoint of 120°F. Weather data came from the U.S. city of

TABLE 1. Percent make up of water heater units.

Bedrooms	Census % [17]	Model %	Modeled Units	Gallons
0	2.4%	0%	0	40
1	10.9%	13.3%	1320	40
2	26.3%	26.3%	2630	40
3	39.6%	39.6%	3960	50
4	16.5%	16.5%	1650	80
5*	4.4%	4.4%	440	80

\*The U.S. Census Bureau records this as 5+ bedrooms.

Portland, Oregon. The number of bedrooms varied to match the census data. The event generator was a macro-enabled Excel spreadsheet, with an event generating macro written in Visual Basic that generates an annual draw schedule for a single set of inputs; a typical annual schedule has around 15,000 water draws per household. For this work, the macro was modified to produce annual water draw schedules for 10,000 households.

The EWHs themselves were modeled using the characteristics of smart-grid enabled units. These units use CTA-2045 application-layer messages to communicate with a DERMS. Three sizes of EWH were used: 40-, 50-, and 80-gallon. The number of bedrooms in the household influences the size of the EWH. One and two bedroom households were assigned 40-gallon EWHs, three bedroom households 50-gallon EWHs, and four and five bedroom households 80-gallon EWHs. The details of the 10,000 unit aggregation are shown in Table 1.

The CTA-2045 standard defines a universal communication module (UCM), which is a standardized appliance socket that facilitates communication between the appliance and a grid operator or aggregator [13]. The standard is being adopted by both manufacturers and jurisdictions. The State of Washington codified CTA-2045 in 2019; the law will require all EWHs manufactured after 2020 to be CTA-2045-compliant as a condition for sale and installation within the state [17].

By using the CTA-2045 standard, an aggregator can read water heater properties such as tank size, element power, and *energy take*, the latter being explain in detail in Section II.B. An aggregator can also issues several different commands to the water heaters, of which the *shed* and *load up* commands are relied upon for this work [13]. These commands are more dynamic than the on/off switching utilized by the Great River Energy aggregation [11]. The commands are service-oriented in that they request services from the appliance, the response to which is left to the interpretation of the manufacturer. This is in contrast to the direct-control approach, wherein a grid operator has the freedom to adjust customer setpoints, as in the Kapsalis and Loukas study [7].

### B. EWH MODELING

The EWHs heat water to a setpoint, which for this work is 120°F for all units. The tank is insulated, but thermal energy slowly dissipates from the tank. When hot water is drawn from the tank, it is replaced by cold water sourced from the

water utility. When the tank temperature drops below a deadband temperature, heater elements turn on to increase water temperature back to the set point. This represents normal operation for an uncontrolled water heater [18]. The model assumes that the temperature of the inlet cold water is  $51^{\circ}F$  based on the average cold water temperature in Portland, Oregon. Heuristically, it was found that a 50 gallon EWH lost 36.3 Wh to ambient conditions of  $70^{\circ}F$  over the course of an hour. The surface area for this measured 50 gallon EWH was used to scale ambient losses for 40 gallon and 80 gallon units.

A smart-grid enabled EWH can be operated as a controlled load. The unit can be given the CTA-2045 *shed* command to tell the water heater to lower its minimum temperature setpoint. The EWH will follow this command past its normal deadband temperature, the temperature at which the unit would ordinarily begin reheating. The EWH will continue to follow this *shed* command until the temperature in the unit drops below the manufacturer's low-limit comfort shutoff temperature. When an EWH drops below this low-limit comfort shutoff temperature, the EWH asserts local control and stops responding to the *shed* command from the grid operator. The EWH starts reheating and continues to reheat until reaching the deadband temperature, at which point the EWH cedes control back to the grid operator.

By continuously sending *shed* commands to the EWH aggregation, the EWH temperatures will oscillate between their deadband temperature and their low-limit comfort shutoff temperature. For this research, the deadband temperature was set to  $117^{\circ}F$ , and the low-limit comfort shutoff temperature was set to  $114^{\circ}F$ . These values derive from observations in other work [19]. By operating the aggregate EWH load in this manner, the EWH aggregation can provide a steady load available for dispatch.

Each EWH has two resistive heating elements. For the smart-grid enabled water heaters modeled in this work, these elements are rated at 4500 W. Only a single coil may operate at a time. The modeled EWHs draw the entire 4500 W of power when a heating element is on.

The *energy take* of an EWH describes the amount of electrical energy an EWH can import. As a water heater loses thermal energy, the tank temperature decreases. Thermal losses occur parasitically due to radiative, conductive, and convective cooling to the environment. Losses also occur due to customer water draws, in which case cold inlet water displaces the withdrawn hot water. In either case, loss of thermal energy results in an increase in *energy take*. Whenever the measured tank temperature is below the setpoint temperature, the EWH has a positive *energy take*, which is a measure of the electrical energy that can be imported from the grid to raise the temperature back to the setpoint. As such, water heaters provide a means for *incrementing load* within an electrical balancing area, which is effectively equivalent to *decrementing generation*.

*Energy take* is measured in Wh, and is defined as the energy necessary to heat the volume of water in the storage tank to

the setpoint temperature.

$$Q = m c (T_{setpoint} - T_{tank})$$

Here,  $m$  is the mass of the water in the tank, and  $c$  is the specific heat of that water,  $4180 \text{ kJ}/(\text{kg } ^{\circ}C)$ . This assumes 100% efficiency from the resistive heating coils. The *energy take* is zero when the temperature in the tank is equal to the setpoint temperature.

### C. EWH DISPATCH

EWHs respond to dispatch commands from the aggregator. For this model, the aggregator is only able to control EWHs that are responding to the *shed* command. Those EWHs that have gone past their low-limit comfort shutoff temperature have asserted local control, so they cannot be dispatched by the aggregator.

The aggregator makes a decrement bid into an energy market. Typically, a decrement bid is a commitment by a generator to reduce its generation for a given period of time. However, an aggregator "reduces generation" indirectly by increasing load, in this case by turning on a large aggregation of EWHs. Supposing the aggregator's bid clears the market, the aggregator will dispatch EWHs to cover this bid. The dispatch algorithm sorts through the available EWHs and calls on some of them to match the bid over each 5 minute operating interval. The aggregator turns on these individual EWHs by sending them the CTA-2045 *load up* command.

A five minute dispatch window is used because it matches the fastest time intervals over which real-time energy markets operate. A dispatch algorithm was written that sorts EWHs by *energy take* from highest to lowest at each 5 minute interval. The dispatch algorithm selects the units with the highest *energy take* first that are under control of the aggregator (i.e. not heating by local control). It then calculates the amount of energy owed over the course of each 5 minute interval and dispatches EWHs to meet that load.

As the bid period continues, EWHs that are heating may reach their setpoint temperature. At this point, the EWH stops heating and becomes idle. At each successive interval over which the bid lasts, the aggregator continues to dispatch units that were already heating and still able to heat. The aggregator assigns additional EWHs to cover for the EWHs that have stopped heating.

To properly account for the utility's energy commitments, the authors assumed that water heaters are part of the demand forecast. Water heaters that have already started heating due to local control are not allowed to cover the energy required by the decrement bid. Additionally, if a dispatched EWH would have turned on in a later 5 minute interval of the bid period, the energy required to heat the EWH from its low-limit comfort level to its deadband would be added to the energy requirement of those later 5 minute intervals. At the end of the bid period, all water heaters return to either local control if the water heaters are below their low-limit comfort temperatures or to their idle state where they can receive and respond to *shed* commands.

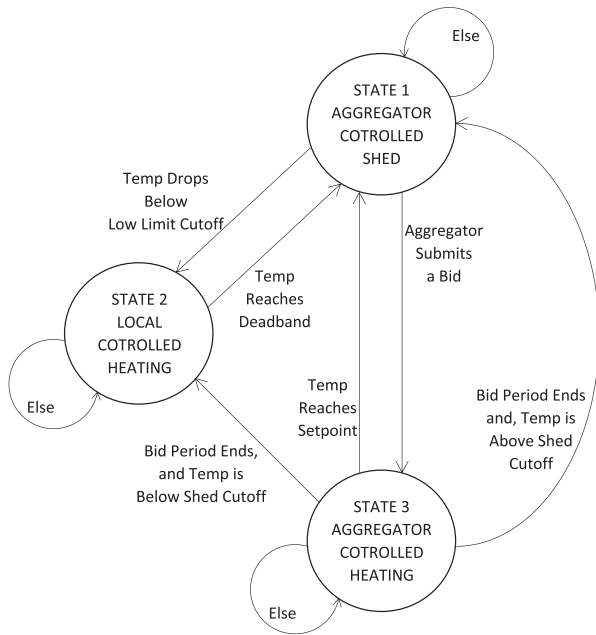


FIGURE 1. State diagram for electric water heater model.

**D. SIMULATING THE EWH AGGREGATION**

The 10,000 water draw profiles were loaded into an array within Matlab. EWHs were assigned an initial *energy take* and state (shedding or heating) through a random number generator. The EWHs were first simulated to develop a base case that did not include any bidding, thereby providing a reference aggregate *energy take* profile to which other experiments could be compared.

*Energy take* is affected by water draws, ambient thermal losses, and locally-controlled heating. The EWHs cycle between two states, as shown in Figure 1. In State 1, *Aggregator-Controlled Shed*, EWHs continuously received the *shed* command, gaining *energy take* as they lose thermal energy to ambient conditions and water draws. In State 2, *Locally-Controlled Heating*, they gain thermal energy due to internal control while also losing thermal energy to ambient conditions and water draws.

A full year of data were simulated using the algorithms in Table 2. Every five minutes, the *energy take* and state of each individual EWH are stored in an array for all 10,000 units. The new *energy take* values and states, along with the water draw schedules, are used to calculate the next *energy take* values and states of the EWHs.

Participation in aggregated dispatch does not result in customers changing their hot water draw profiles. An EWH that is dispatched by an aggregator defers reheating to a later time, which is a temporal shift in energy consumption rather than a change in energy consumption magnitude. Customer water draw profiles are expected to stay unchanged so long as the water heater temperature does not drop below the minimum threshold temperature. As such, customers are unlikely to notice that their water heaters are participating in the programs. Studies by both Pacific Northwest

TABLE 2. EWH calculations at each state.

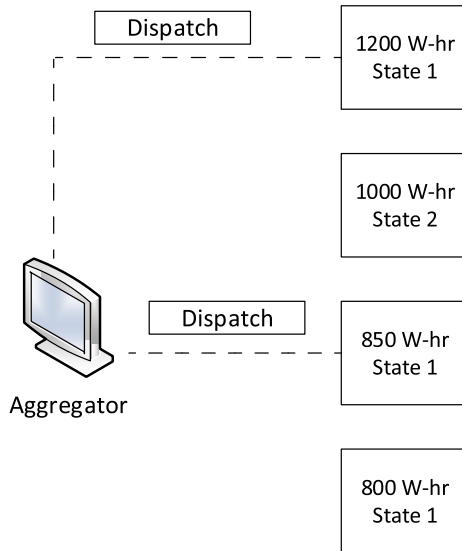
State 1, Aggregator-Controlled Shed	$Q_k = Q_{k-1} + Q_{amb} + Q_{k,draw}$ if $Q_k < Q_{LowLimit}$ Set State = 1 else Account for Heating Set State = 2
State 2, Local-Controlled Heating	$Q_k = Q_{k-1} + Q_{amb} + Q_{k,draw} - Q_{heat}$ if $Q_k > Q_{db}$ Set State = 2 else Account for Shedding Set State = 1
State 3, Aggregator-Controlled Heating	$Q_k = Q_{k-1} + Q_{amb} + Q_{k,draw} - Q_{heat}$ if $Q_k > Q_{setpoint}$ Set State = 3 else Account for Shedding Set State = 1

National Laboratory and Bonneville Power Administration have shown that such modifications to tank temperatures do not significantly impact customer participation satisfaction [12], [20].

An algorithm to dispatch EWHs to cover cleared bids was scripted using Matlab. This algorithm established a State 3, *Aggregator-Controlled Heating*. The transitions between all three states are illustrated in Figure 1. When a dispatch is activated, an array is built to describe the aggregation’s net *energy take* and the states of all EWHs for the current day and the next day. The base case is used to describe the energy states until the dispatch occurs. Once the dispatch occurs, the aggregation is resimulated.

The EWHs are sorted from high *energy take* to low *energy take*. The unit with the highest *energy take* is assigned to cover the bid if the unit is in State 1. To assign a unit to a bid means changing the unit’s state from State 1 to State 3. If the EWH is in State 2, it is skipped because it is not responding to aggregator commands. This process is illustrated in Figure 2. After a unit is assigned to cover the bid, the energy that unit can provide across the next five minute period is calculated and subtracted from the total energy required across the current five minute interval. This process of assigning the highest *energy take* units continues until the energy required to cover the bid is met for the current five minute interval. At this point the units are simulated across the five minute interval using the algorithms from Table 2. The ending *energy take* values and states are then recorded.

After resimulating over a five minute interval, the program checks to see if the bid time period has ended. If the bid has ended, the program reassigns units that were in State 3 back to State 1. If the bid has not ended, the program lets all the units in State 3 remain there, and again calculates the necessary



**FIGURE 2.** The aggregator dispatches the units with the highest energy take, skipping those in State 2.

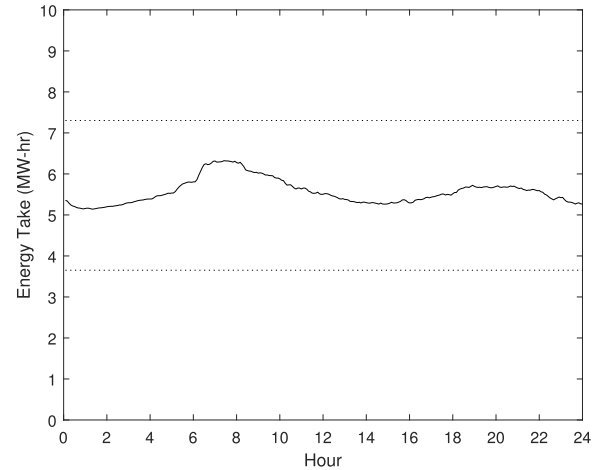
energy to provide across the next five minute interval. The energy provided by the units in State 3 is calculated across the new five minute interval and subtracted from the required energy to cover the bid across that interval. If additional units are required, EWHs are dispatched according to the dispatch algorithm, with the highest energy take units being selected first. The simulation repeats across this interval and continues until the bid period ends. Once the bid period ends, all units are restored to either States 1 or 2.

### III. RESULTS OF SIMULATING AGGREGATED EWH

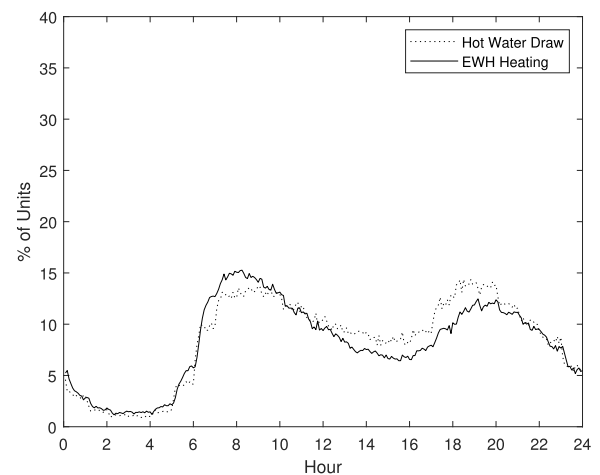
Simulations were used to test four cases: no bidding, a one hour bid, an 18 hour bid, and a 5 minute bid. For consistency, all of the test bids start at 02:00 on the 6th of June. Bid power and duration were chosen in order to demonstrate the impact that bids have on the aggregate system. Bids were also chosen that stressed the system by requesting either more energy than was available from the aggregate or more power than was available. By analyzing the results from these simulated bids, one may observe how water heater units are allocated to meet the bid, how the aggregate energy take changes as a result of a bid, and how the aggregate energy take recovers once the bid is complete. Bids were not selected by considering the perspective of a dispatcher, wherein a bid size needs to be determined to meet demand requirements. Rather, bid sizes were chosen to perturb the system and observe the results.

#### A. NO-BIDDING CASE

In this case, no bids were made, so the simulated data only relied on the water heaters to act due to ambient heat losses and water draws. Figure 3 shows the aggregation's energy take over the course of the day. The upper horizontal line shows the energy take if all units were at their low-level comfort shutoff temperature. This is the point at which the



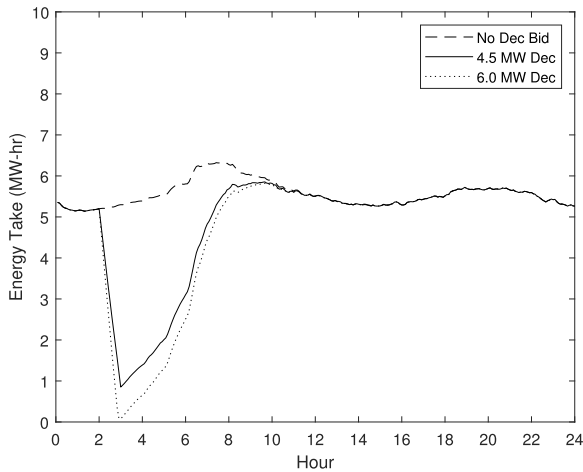
**FIGURE 3.** The simulated energy take of the EWH aggregation with no bids.



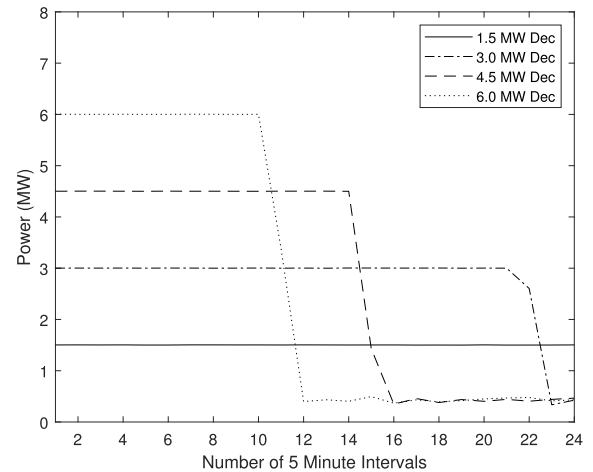
**FIGURE 4.** The percent of the EWH aggregation heating and percent of households experiencing hot water draws during a summer weekday.

EWHs assert local control and start heating themselves. The lower horizontal line shows the energy take if all the water heaters were at their shed deadband temperature. With no bids, this is the minimum energy take the EWHs would have when responding to an aggregator's constant shed commands. The amount of energy take is a key metric in that it shows how much energy the EWH aggregation can draw from the grid.

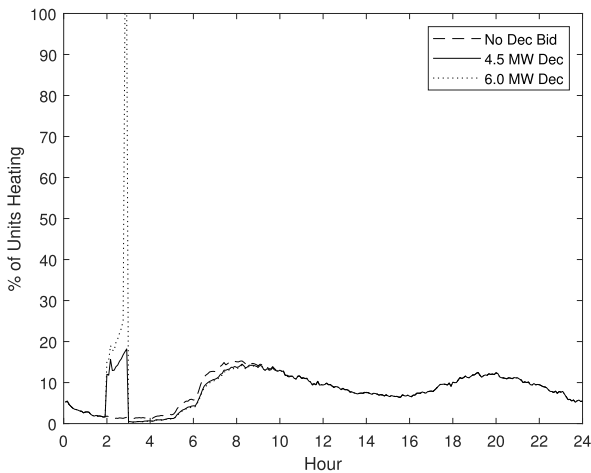
The EWHs see a rise in energy take in the morning hours around 06:00 to 08:00, and a rise in the evening from 18:00 to 22:00. The main cause of variation here comes from water heaters responding to water draws. Figure 4 shows the percent of water heaters in the aggregation experiencing a water draw, and the percent of EWHs in the aggregation that are heating. The EWHs heating appears to slightly lag the hot water draws, as expected. During the morning peak, a little over 15% of EWHs are heating. The percent of EWHs heating is another key metric as it speaks to the power capacity of the aggregate.



**FIGURE 5.** The energy take of the EWH aggregation responding to a 1 hour bid from 02:00 to 03:00.



**FIGURE 7.** The power absorption of the EWH aggregation responding to a 2 hour bid from 02:00 to 04:00.



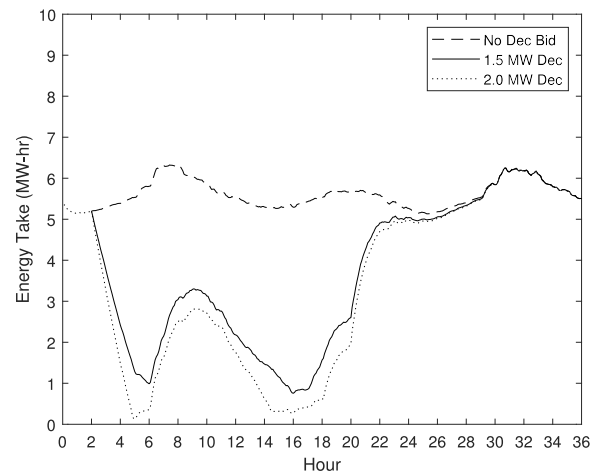
**FIGURE 6.** The percent of the EWH aggregation heating while responding to a 1 hour bid from 02:00 to 03:00.

**B. ONE HOUR BIDDING CASE**

As Figure 3 shows, for an aggregation of 10,000 EWHs, a little over 5 MWh of energy take is reliably available throughout the day. This bidding simulation examines one hour bids of 4.5 MW and 6 MW to analyze the effects of a bid that meets its obligations versus an over bid, which fails to meet the bid commitment.

Figure 5 shows a nearly-linear drop in energy take across the hours of 02:00 to 03:00 as EWHs are dispatched and start heating. The 4.5 MW bid drops the energy take to around 1 MWh while the 6 MW bid drops the energy take to 0 MWh. Figure 6 shows the percent of EWHs heating. In the 4.5 MW bid, about 20 to 25 percent of the EWHs are dispatched during the bid period. In the 6 MW bid, the dispatch looks similar to the 4.5 MW bid until the last two intervals when it dispatches 100% of its units but still cannot cover the bid. After the bid period, the EWH aggregation starts to recover. A 2% settling time compared to the no-bid case occurs around 10:00 am.

Power output of the system is shown in Figure 7 for bids over a two hour period. The aggregation offers power until



**FIGURE 8.** The energy take of the EWH aggregation responding to an 18 hour bid from 02:00 to 20:00.

it runs out of energy take and power drops to near zero. Due to the consistent presence of hot water draws and ambient losses, the floor of the power output remains around 0.4 MW.

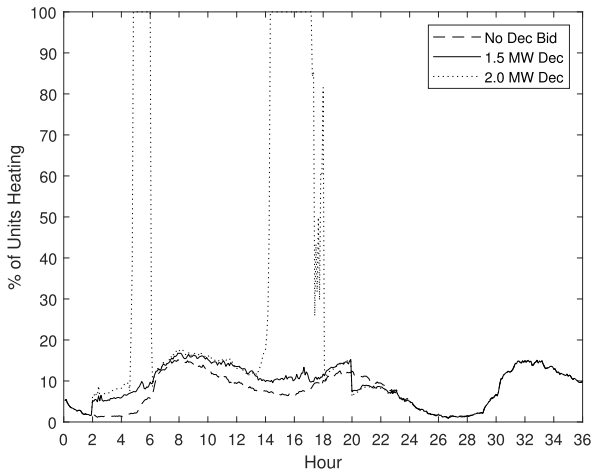
**C. 18 HOUR BIDDING CASE**

An 18 hour bid offers a much smaller amount of power over a long period of time. 1.5 MW and 2.0 MW bids were simulated over a time period from 02:00 to 20:00. Figure 8 shows the energy take throughout the bid period starting at midnight on the day of the bid and going to noon on the following day.

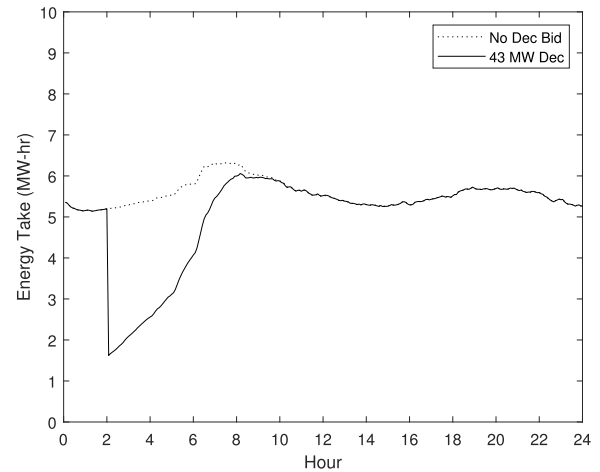
Figure 9 shows the percent of the EWHs heating over each interval. The 1.5 MW bid has no problems absorbing power, but the 2.0 MW has two periods from 04:30 to 06:00 and 14:00 to 18:00 over which the aggregation fails to absorb the necessary power despite all the EWHs heating. Figure 10 shows the aggregation’s power absorption throughout the bid period.

**D. 5 MINUTE BIDDING CASE**

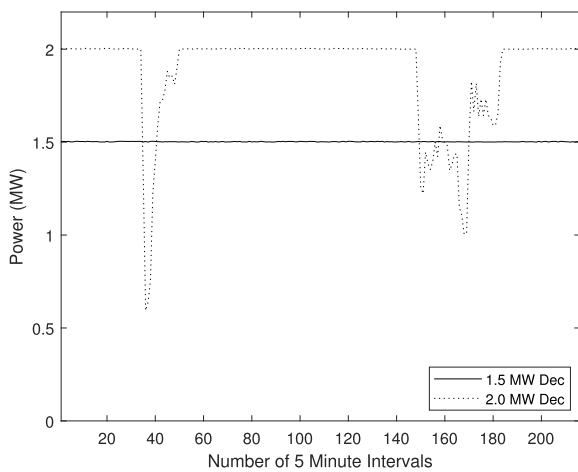
Previous bids were limited by the EWH aggregation running out of energy take; the bids were limited by the aggregation’s



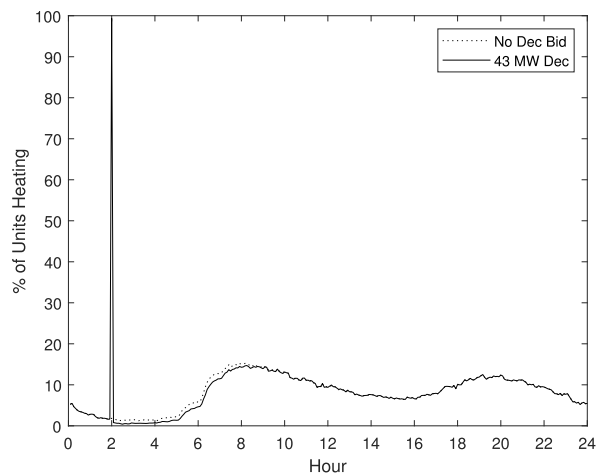
**FIGURE 9.** The percent of the EWH aggregation heating while responding to an 18 hour bid from 02:00 to 20:00.



**FIGURE 11.** The energy take of the EWH aggregation responding to a 5 minute bid from 02:00 to 02:05.



**FIGURE 10.** The power absorption of the EWH aggregation responding to an 18 hour bid from 02:00 to 20:00.



**FIGURE 12.** The percent of the EWH aggregation heating while responding to a 5 minute bid from 02:00 to 02:05.

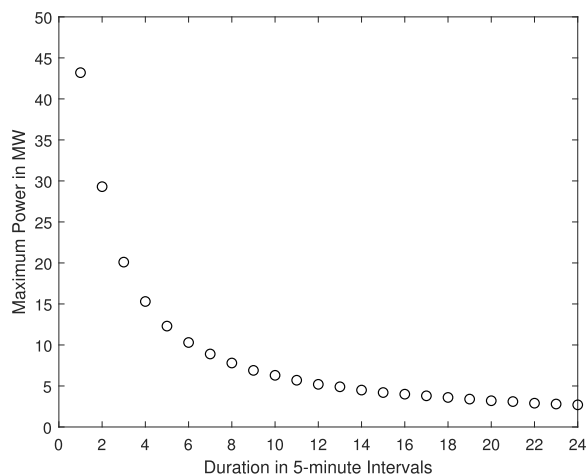
ability to store energy. Shorter period bids can be limited by power rather than *energy take*. Figure 3 shows roughly 5 MWh of *energy take*, which would equate to 60 MW of power over a single 5 minute period. However, the EWHs themselves only have a 4500 W coil, so the aggregation of 10,000 units can absorb a maximum of 45 MW.

Additionally, the unit dispatch can only call on EWHs in State 1, which are the units under control of the aggregator. When units have dropped below their low-limit cutoff, State 2, they cannot be dispatched. Figure 11 shows the results of a 45 MW bid simulated across a 5 minute period. A sizable amount of energy is absorbed across the 5 minute period, yet about 1.7 MWh of *energy take* remain. Figure 12 shows the amount of units heating, and it appears 100% of the units are heating. Checking the simulation’s numerical output, 99.7% of the units are actually heating and 0.3% remain shedding. With only 43 MW being absorbed from the aggregation, this bid situation has a surplus of *energy take*, but is limited by the aggregated power of the heating elements.

#### IV. DISCUSSION

EWHs have a primary purpose: to provide hot water to customers. An EWH aggregation could provide a grid service, but this should be considered secondary to their primary purpose. Otherwise, customers would likely be dissuaded from participating in a DERMS program. This model assumes that the EWHs are able to provide hot water at all times, and that any grid service functionality comes second.

The simulations of the 10,000 EWH aggregation demonstrate the ability of an aggregator to follow decrement bids in a 5 minute market. The model consisted of 39.5% of households with 40 gallon units, 39.6% with 50 gallon units, and 20.9% with 80 gallon units. This could be a particularly useful model for the State of Washington, which recently passed a law requiring new EWHs be CTA-2045 compliant [17]. The model also assumes assets that had turned on before the bid or would turn on due to ambient heat losses could not be used as part of the aggregation, as these assets should be considered part of the load forecast. The model



**FIGURE 13.** The maximum bid the EWH aggregation can cover over a given number of five minute intervals starting at 02:00.

additionally assumed that the power availability of the units was 4500 Watts, which is not an unreasonable assumption as most water heaters have power ratings of this value.

With these assumptions in mind, a 10,000 unit EWH aggregation would have around 5 MWh of energy available at any given time. The available power would be related to the units that were not heating or forecasted to be heating. The maximum available power would be the sum of the power of all heating elements in all of the units, or 45 MW. However, since some of these would already be heating (State 2), this could be de-rated by as much as 20%, to around 36 MW.

When overbidding occurs, the aggregation is unable to deliver the necessary power over the required time period. This leads to the power drawn by the aggregation diminishing in the later intervals. For short, high power bids, the system does not recover, and the power diminishes to about 0.4 MW, or the continuous energy coming from the hot water draws as in Figure 7. However, when fewer megawatts are bid across longer intervals, the hot water draws could allow the system to recover back to full power, except for a few non-compliant 5 minute intervals, as shown in Figure 10.

Because overbidding would result in insufficient power over certain intervals, the dispatch algorithm should assess the bid prior to dispatch. The demonstrated algorithm dispatches units to cover the bid energy until the bid energy can no longer be covered and the system starts to fail. An alternative approach would be to recognize first that the aggregator cannot cover the bid, and then even out deficiencies in the bid over the entire time interval. For short bids, like a 1 hour, 6 MW bid, the solution is to simply offer 5.2 MW over the entire hour. The dispatch algorithm can be simulated across any number of 5 minute intervals to determine how much power can be bid without failing to deliver across any individual interval. Figure 13 illustrates the maximum power that can be bid across any number of five minute intervals up to two hours, yet this tool cannot be utilized at bid time without perfect forecasting. A probabilistic model could be developed to account for imperfections in forecasting. For

longer intervals like the 18 hour period, a 2 MW bid that reduces the amount of disturbance across all intervals would be more desirable than a dispatch that assigns units until failure. Again, this solution also depends on a tool utilizing a probabilistic forecast.

After nearly depleting themselves, the aggregation of EWHs takes roughly 7 hours to recover back to the non-bid *energy take* level. Here there are two settling times, the energy settling time for when *energy take* recovers to within 2% of the non-bid levels and the power settling time for when the number of units heating returns to within 2% of the non-bid levels. After these settling times, the aggregate system effectively returns to its normal energy state as if no bids had occurred.

## V. CONCLUSION

Standards bodies within the electric power industry have recently developed several open protocols designed to facilitate the aggregation and dispatch of large numbers of residential distributed energy resources such as EWHs. These include CTA-2045, the Modular Communications Interface for Energy Management; IEEE 2030.5, the Smart Energy Profile Application Protocol; and, OpenADR, an open protocol for automated demand response. Consequently, utilities, aggregators, and software developers are now developing software systems and utility dispatch services capable of aggregating and dispatching large numbers of residential-scale distributed energy resources.

This study considered the aggregation of EWHs enabled with CTA-2045. This smart-grid protocol allows an aggregator to dispatch EWHs *en masse* in order to provide decrement capacity, while concurrently allowing EWHs to provide hot water as their primary service. This study provides insights that should prove useful to the electric power industry. First, this study demonstrated the energy and power capacities that are available when aggregating large numbers of EWHs, and how those capacities change over time. An aggregation of 10,000 EWHs has at least 5 MWh of *energy take* available and can dispatch between 30 and 43 MW at any given time; these numbers are scalable to larger aggregations. Second, the study showed that distributed energy resources, which have finite energy needs, have limited *energy take* available to contribute to energy bids, and that they can take multiple hours to recover to pre-bid *energy take* states. And third, the study showed that dispatch of distributed energy resources is limited by the aggregate power capacities of the assets minus those that are not available because they are below their low-limit comfort shutoff temperature.

This research simulated the dispatch of a large number of EWHs, each featuring a unique hot water draw profile. The work demonstrated the ability of aggregations of EWHs to provide “generator”-like decrement capacity to a grid operator through various-sized energy bids. For the purpose of making bids, an aggregator would want to predict an amount of power and dispatch duration that could be offered based on forecast data, and to understand how the *energy take*



capacity would rebound after the bid. Doing so would require a probabilistic model, and should be an endeavor of future work.

## REFERENCES

- [1] J. Giri, "Proactive management of the future grid," *IEEE Power Energy Technol. Syst. J.*, vol. 2, no. 2, pp. 43–52, Jun. 2015.
- [2] J. Kondoh, N. Lu, and D. J. Hammerstrom, "An evaluation of the water heater load potential for providing regulation service," *IEEE Trans. Power Syst.*, vol. 26, no. 3, pp. 1309–1316, Aug. 2011.
- [3] J. Laurent and R. Malhame, "A physically-based computer model of aggregate electric water heating loads," *IEEE Trans. Power Syst.*, vol. 9, no. 3, pp. 1209–1217, Aug. 1994.
- [4] N. Fitzgerald, A. M. Foley, and E. Mckeogh, "Integrating wind power using intelligent electric water heating," *Energy*, vol. 48, no. 1, pp. 135–143, Dec. 2012.
- [5] M. Roux, M. Apperley, and M. Booyens, "Comfort, peak load and energy: Centralised control of water heaters for demand-driven prioritisation," *Energy Sustain. Develop.*, vol. 44, pp. 78–86, Jun. 2018.
- [6] J. Li, Z. Wu, S. Zhou, H. Fu, and X.-P. Zhang, "Aggregator service for PV and battery energy storage systems of residential building," *CSEE J. Power Energy Syst.*, vol. 1, no. 4, pp. 3–11, 2018.
- [7] V. Kapsalis and L. Hadellis, "Optimal operation scheduling of electric water heaters under dynamic pricing," *Sustain. Cities Soc.*, vol. 31, pp. 109–121, May 2017.
- [8] V. Kapsalis, G. Safouri, and L. Hadellis, "Cost/comfort-oriented optimization algorithm for operation scheduling of electric water heaters under dynamic pricing," *J. Cleaner Prod.*, vol. 198, pp. 1053–1065, Oct. 2018.
- [9] M. Toro, A. Samra, and S. Sim, "Evaluation of load management program in South Florida," in *Proc. SOUTHEASTCON*, Dec. 2002, pp. 225–227.
- [10] *Actively Adapting to the Changing Electricity Sector*, document SB 978, Oregon Public Utility Commission, Sep. 2018.
- [11] R. Hledik, J. Chang, and R. Lueken, "The hidden battery: Opportunities in electric water heating," Brattle Group, Washington, DC, USA, Tech. Rep., 2015. [Online]. Available: [https://rpsc.energy.gov/tech-solutions/sites/default/files/attachments/Brattle-Group\\_The-Hidden-Battery-01-25-2016.pdf](https://rpsc.energy.gov/tech-solutions/sites/default/files/attachments/Brattle-Group_The-Hidden-Battery-01-25-2016.pdf)
- [12] D. Hammerstrom *et al.*, "Pacific Northwest GridWise Testbed demonstration projects; Part I. Olympic peninsula project," Pacific Northwest Nat. Lab., Richland, WA, USA, Tech. Rep. PNNL-17167, Oct. 2007.
- [13] *Modular Communications Interface for Energy Management*, Standard ANSI/CTA-2045, Consum. Electron. Assoc., Feb. 2010.
- [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] B. Hendron, J. Robert, and G. Barker, "Tool for generating realistic residential hot water event schedules," in *Proc. SimBuild Conf.*, New York, NY, USA, Aug. 2010, pp. 1–11.
- [16] *U.S. Residential Household Data*, U.S. Bureau Census, Amer. Fact Finder, Suitland, MD, USA, 2018.
- [17] *Concerning Appliance Efficiency Standards*, document Second Substitute House Bill 1444, Washington State Legislature, Jul. 2019.
- [18] M. Pipattanasomporn, M. Kuzlu, S. Rahman, and Y. Teklu, "Load profiles of selected major household appliances and their demand response opportunities," *IEEE Trans. Smart Grid*, vol. 5, no. 2, pp. 742–750, Mar. 2014.
- [19] T. Clarke, "Aggregation of electric water heaters for peak shifting and frequency response services," M.S. thesis, Portland State Univ., Portland, OR, USA, Jun. 2019.
- [20] "Water heater demonstration report," Bonneville Power Admin. Tech. Innov. Project 336, Portland, OR, USA, Tech. Rep. CTA-2045, Nov. 2018, [Online]. Available: <https://www.bpa.gov/EE/Technology/demand-response/Documents/Demand%20Response%20-%20FINAL%20REPORT%20110918.pdf>

**KEVIN MARNELL** (Member, IEEE) received the B.S. degree in chemical engineering from Oregon State University, Corvallis, OR, USA, in 2004, and the M.S. degree in electrical engineering from Portland State University, Portland, OR, USA, in 2019. He currently works as a Protections and Controls Engineer with PacifiCorp, Portland.

**CONRAD EUSTIS** received the Ph.D. degree in engineering and public policy from Carnegie Mellon University, in 1985. He was the Director of Retail Technology Strategy with the Portland General Electric, until retiring in August 2019. He worked on advanced energy systems for 42 years particularly on development of technology that enabled energy delivery with 100% sustainable sources.

**ROBERT B. BASS** (Member, IEEE) received the Ph.D. degree in electrical engineering from the University of Virginia, Charlottesville, VA, USA, in 2004. He is currently an Associate Professor with the School of Electrical and Computer Engineering, Portland State University, Portland, OR, USA. His research interests pertain to electric power distribution assets, power quality, renewables integration, and distributed energy storage systems.