



Pacific Northwest
NATIONAL LABORATORY

*Proudly Operated by **Battelle** Since 1965*

Loads as a Resource

Frequency Responsive Demand

November 2015

K Kalsi
T Williams
LD Marinovici
M Elizondo
J Lian

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.** Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY

operated by

BATTELLE

for the

UNITED STATES DEPARTMENT OF ENERGY

*under Contract DE-AC05-76RL01830***Printed in the United States of America**

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <<http://www.ntis.gov/about/form.aspx>>
Online ordering: <http://www.ntis.gov>



This document was printed on recycled paper.

(8/2010)

Frequency Responsive Demand

K Kalsi
T Williams
LD Marinovici
M Elizondo
J Lian

November 2015

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Executive Summary

With large-scale plans to integrate renewable generation driven mainly by state-level renewable portfolio requirements, more resources will be needed to compensate for the uncertainty and variability associated with intermittent generation resources. Distributed assets can be used to mitigate the concerns associated with renewable energy resources and to keep costs down. Under such conditions, performing primary frequency control using only supply-side resources becomes not only prohibitively expensive, but also technically difficult. It is therefore important to explore how a sufficient proportion of the loads could assume a routine role in primary frequency control to maintain the stability of the system at an acceptable cost.

In FY13 (see [1]), a demand-side primary frequency strategy was proposed to regulate the frequency of the system to its nominal value and restore power balance for a multi-machine power system model. The load control law was designed in a hierarchically decentralized manner consisting of two interactive decision layers. In the first layer, a supervisory controller is responsible to gather system-level information (e.g., power flow, system topology, generation and load forecast, available responsive loads, among others), and determine the optimal gains for the responsive loads on each bus every few (e.g., 15-30) minutes. In the second layer, each decentralized load switched ON/OFF probabilistically in real time based on local frequency measurement so that the aggregated load response under each bus matches the desired power determined by the first layer. The control gain design in the first layer was based on the decentralized robust control theory, while the load switching probabilities were designed using Markov chains. The proposed decentralized robust control strategy was tested as a proof-of-concept using IEEE test systems in the MATLAB Power Systems Toolbox (PST).

In FY14 (see [2]), the controller design proposed in FY13 was extended to deal with more realistic power system models. Furthermore, the framework was extended to include network-preserving power system models. This allowed the design to fully respect the topology of the power system network, which is crucial for primary frequency control problems. Such an extension is highly nontrivial as the load bus transient frequency dynamics and their interactions with the generation side dynamics have not been well studied in the literature. A rigorous theoretical analysis was performed to prove the closed-loop system stability under demand-side frequency control. Specifically, it was shown that with only local frequency measurements, the frequency synchronization and asymptotic stability around post-fault equilibrium can still be satisfied as long as the control gains were chosen to be non-negative. The modified control strategy was implemented on the Western Electricity Coordinating Council (WECC) model using PowerWorld. Simulation studies and performance metrics were designed to analyze the impacts of large-scale deployment of frequency responsive load. The results indicated that the proposed frequency-based load control strategy showed improved frequency recovery in terms of steady state error and maximum frequency deviation. The simulations showed that the response of intertie power flows depended on relative location of disturbance and responsive load and generation. An important aspect of the enhanced frequency-based load control strategy was that the control gain can be adapted, by the system operator, based on the available controllable load resources and changing operating conditions.

In FY15, hardware testing and validation of the frequency-based load control strategy were performed. Utilizing two different types of water heaters under realistic use conditions, the response of the appliance to the control signals was validated and improved. These improvements were incorporated

into the appliance models used in PowerWorld for the rest of the FY15 work, which were used to perform large-scale population studies on the WECC system.

Utilizing the hardware-validated models, system-wide studies were conducted to investigate the appropriate selection of control gain for the demand-side frequency controller developed in FY14. In particular, the use of the controller is considered to help meet the requirements of the North American Electric Reliability Corporation (NERC) BAL-003-1 Frequency Response and Frequency Bias Setting Reliability Standard [3]. The simulation results demonstrated the frequency-based load control strategy could provide benefits to the overall grid in meeting the BAL-003-1 requirement. The simulations also provided insights into the ability of existing appliance designs to provide desirable performance at aggregate level, without customer dissatisfaction or actions that may potentially cause impact on product life.

Acknowledgments

The authors are grateful for the comments and feedback provided by Jason Fuller, Jeff Dagle and Frank Tuffner at the Pacific Northwest National Laboratory, Joe Eto at Lawrence Berkley National lab, and Phil Overholt at the Department of Energy.

Acronyms and Abbreviations

NERC	North American Electric Reliability Corporation
TCL	Thermostatically Controlled Loads
HVAC	Heating, Ventilation, and Air Conditioning
PST	Power System Toolbox
PNNL	Pacific Northwest National Laboratory
HPWH	Heat Pump Water Heater
ERWH	Electric Resistance Water Heater
WECC	Western Electricity Coordinating Council
UDM	User-Defined Model
IFRO	Interconnection Frequency Response Obligation
FRO	Frequency Response Obligation
BPA	Bonneville Power Administration

Table of Contents

Executive Summary	iii
Acknowledgments.....	v
Acronyms and Abbreviations	vii
1.0 Introduction	4
2.0 Hierarchical demand-side frequency control strategy	6
2.1 Overview of control structure	6
2.2 Design of feedback control gain to meet NERC reliability standard	7
2.3 Device-level control model	8
3.0 Study of system-wide impacts of large-scale deployment of frequency-responsive load control.....	9
3.1 Modeling hierarchical demand-side frequency control in PowerWorld	10
3.2 WECC test scenarios.....	11
3.3 Simulation results.....	13
3.3.1 Analysis of results for system-wide frequency response	14
3.3.2 Analysis of results for area-wide frequency response.....	15
4.0 Hardware testing of proposed frequency-based load control strategies	19
4.1 Hardware setup used for experiments	19
4.2 Hardware testing and validation.....	23
4.2.1 Power Cycle Testing Results.....	23
4.2.2 Water Heater Load Patterns	24
4.2.3 Aggregate Model Validation.....	26
4.3 Population studies based on hardware testing.....	28
5.0 Conclusions	31
6.0 Future Work.....	32
References.....	33

Figures

Figure 1. Hierarchical structure controller with supervisory level and device layer	6
Figure 2. The 2-state Markov-Chain model of an electrical water heater.....	8
Figure 3. The 4-state Markov-Chain model of a heat pump water heater or HVAC system.....	9
Figure 4 Diagram of user-defined model (UDM) of frequency-responsive load controller in PowerWorld and its interface with the new WECC composite load model [19].....	11
Figure 5. WECC balancing authorities (right) and sub regions for reserve sharing groups (left) [21].....	12
Figure 6. Frequency at a bus in the Northwest area for cases 1 (no controllable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)	14
Figure 7. System overall frequency response for cases 1 (no controlable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)	14
Figure 8. Tie line power flow, north to south, for cases 1 (no controlable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)	15
Figure 9. Contributions of generators to overall frequency response for Cases 1 (no controlable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south).....	15
Figure 10. Contributions of loads to overall frequency response for Cases 1 (no controlable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)	16
Figure 11. Contributions of losses to overall frequency response for Cases 1 (no controlable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)	16
Figure 12. Area overall frequency response of an area in the south for Cases 1 (no controlable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south).....	17
Figure 13. Total area generation of an area in the south for cases 1 (no controlable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south).....	18
Figure 14. Total area load of an area in the south for cases 1 (no controlable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)	18
Figure 12. Hardware configuration in laboratory.	20
Figure 13. Control and monitoring system implementation	21
Figure 14. Water draw and power use over an illustrative 24-hour period.....	22
Figure 15. Power draw of ERWH through transitions of state.	24
Figure 16. Daily Mean Water Heater Usage in the Pacific Northwest [29].....	25
Figure 17. Water heater hourly load shapes in the Pacific Northwest [29].	26
Figure 18. Simulated and observed state populations for ERWH	27
Figure 19. Simulated and observed state populations for HPWH	28
Figure 20. System frequency for varying participation levels of ERWH.	29
Figure 21. System frequency for varying participation levels of HPWH.	29

Figure 22. Primary frequency control performance for ERWH and HPWH as a function of participating device population size.....30

1.0 Introduction

The U.S. electric power grid is traditionally operated by delivering electricity generated by bulk power generation at central plants to the end-users via transmission and distribution networks. In order to ensure reliability of the power grid, peaking power generation is required to provide additional capacity during peak hours, as well as other ancillary services such as load following, contingency reserves, etc. The rapid growth of low-inertia renewable energy resources represents an immense opportunity for the U.S. to minimize its carbon footprint, while presenting a challenge for system operators as traditional “spinning” generation resources are displaced. Furthermore, renewable energy generation is typically non-dispatchable, which substantially increase the need for operational reserves to balance supply and demand instantaneously and continuously [4]-[6]. There is growing recognition that distributed energy resources (DERs – loads, distributed generation, storage, electric vehicles, etc.) represent a great potential to provide synthetic reserve. However, operators have concerns about their controllability and dependability, especially when not under their direct control.

Demand-side frequency control presents a novel and viable way for providing the desired frequency response. Loads can measure frequency locally and change their power consumption after a non-zero frequency deviation is observed, in order to achieve power balance between supply and demand. Several load control implementations have been proposed in [7]-[11]. In all the aforementioned approaches, frequency control was based on overly-simplified power system models (such as a lumped linear system model or a multi-machine model with DC power flow). In addition, device constraints and end-use behavior (such as compressor time delays for thermostatically-controlled loads, and stochastic load arrival and departure) have not been considered in a systematic manner in the controller design. These constraints can significantly influence the aggregated response of the controllable devices, and thus impact the overall frequency dynamics. Furthermore, most previous approaches could not deal with time-varying system conditions, such as power flow across the network, available responsive loads, changing network topology, etc. In FY14, PNNL proposed a hierarchical demand-side frequency control framework ([2]) using a structure preserving power system model, instead of the simplified power system model used in previous studies. Furthermore, in order to respect the time-varying capabilities of end-use loads, the control gains need to be carefully designed by accounting for varying system operating conditions. Extensive simulation studies using different scenarios were performed using the WECC system model in PowerWorld to validate the effectiveness of the newly proposed control design.

In FY15, a coordination strategy is proposed to determine the required controller gains so that the control strategy proposed in FY14 can help meet the North American Electric Reliability Corporation (NERC) BAL-003-1 Frequency Response and Frequency Bias Setting Reliability Standard [3]. The proposed strategy can be used to coordinate the autonomous responses at a system-wide, area-wide and local bus levels. The proposed strategy allows for coordination in a nested way from the system level to area levels, and area levels to bus levels. In a similar fashion, additional levels of coordination could be defined at sub regional levels (between system and area level), and also down to the sub-transmission and distribution levels until reaching the end-use loads. The effectiveness of the proposed coordination strategy will be investigated on the WECC system model using the PowerWorld simulation software package. Detailed end-use load models, such as the WECC composite load model, are implemented in PowerWorld to better approximate the effects of the distribution system as part of the transmission system model. Simulation results will examine how a specific level of response in MW/Hz can be achieved at the system level, while also designing specific levels of MW/Hz contributions at the area levels to help meet

NERC BAL-003-1 standard. This will include investigations into how the penetration level of responsive loads (as a percentage of the total load) impact the effectiveness of the coordinated control.

In addition, hardware testing was also performed in FY15 to verify the applicability of the control logic over a wide range of operating conditions while accounting for unpredictable end-use device behavior and physical device constraints. The testing was carried out on two types of residential-scale water heaters under realistic use conditions. A traditional, resistive heating water heater and a heat-pump water heater were tested with the control strategy developed in FY14. Water heaters have been identified as candidates for frequency control, so validation of this response is necessary for demonstrating the overall benefits and capabilities of the proposed demand-side control scenario. Incorporated into the system-wide WECC studies, the benefits of the frequency-based demand-side control strategy the capabilities of the controller to provide services for meeting both BAL-003-1 frequency requirements, as well as general demand response capabilities, will be demonstrated.

2.0 Hierarchical demand-side frequency control strategy

In this section, first an overview of the hierarchical demand-side frequency control strategy is presented. Next, a systematic method is developed for selecting the load control gain parameters to meet the NERC BAL-003-1 Frequency Response and Frequency Bias Setting Reliability Standard.

2.1 Overview of control structure

In our previous work [2], [12] and [13], a hierarchical control structure was proposed to coordinate controllable loads to provide primary frequency response. The hierarchical structure consisted of both supervisory and device-level control layers, as illustrated in Figure 1. At the supervisory control layer, the optimal feedback gains required from the aggregated load response on each bus are computed to ensure stability of the closed-loop system over a wide-range of operating conditions. The systematic computation of the gains every few minutes (e.g., 10-15 minutes) in the supervisory control layer allows for an adaptation to time-varying system operating conditions. The control gains are then transmitted to the individual devices that react to local frequency measurement based on gain setting received from the supervisory layer. The mode switching probabilities of individual devices are computed based on a Markov-chain model and the measured frequency deviation so that the total population response matches the desired one determined by the supervisory controller. The Markov-chain-based design of local load response rules can fully respect the device constraints (such as temperature setpoint, compressor time delays of Heating, Ventilation, and Air Conditioning (HVAC) units, arrival and departure of the deferrable loads, etc.), which are crucial for ensuring load control programs are deployed in practice.

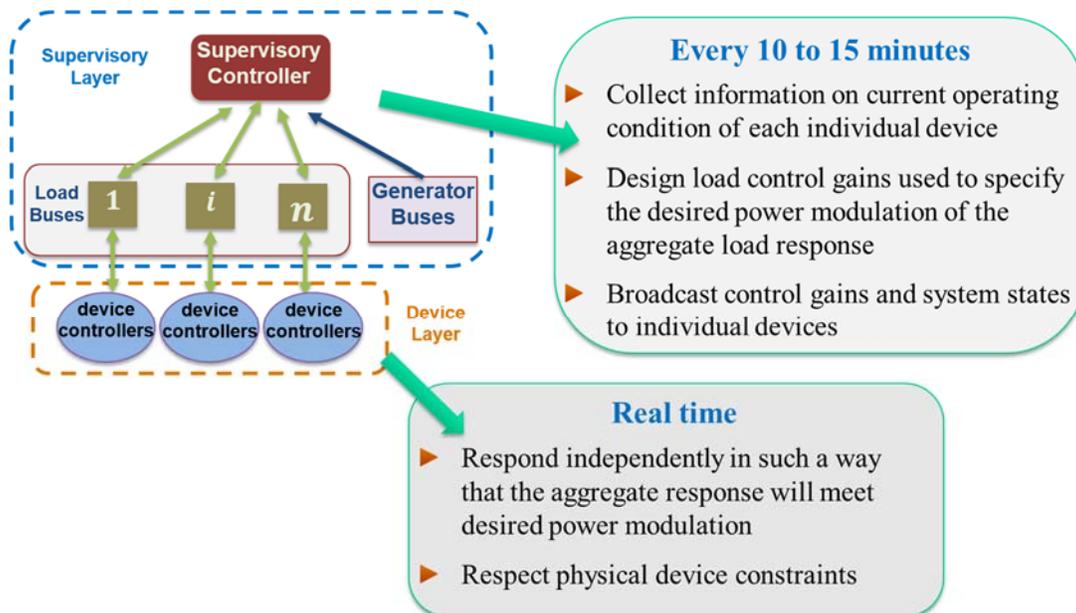


Figure 1. Hierarchical structure controller with supervisory level and device layer

2.2 Design of feedback control gain to meet NERC reliability standard

The NERC BAL-003-1 standard defines requirements to maintain interconnection frequency within acceptable bounds. As discussed in [3], NERC BAL-003-1 states its objective as: “*To require sufficient Frequency Response from the Balancing Authority (BA) to maintain Interconnection Frequency within predefined bounds by arresting frequency deviations and supporting frequency until the frequency is restored to its scheduled value*”. NERC BAL-003-1 establishes system-wide requirements given by the Interconnection Frequency Response Obligation (IFRO), and area-wide requirements given by Frequency Response Obligation (FRO). IFRO is calculated based on contingencies and desired system response, and FRO is calculated as a portion of IFRO, that is being contributed by areas (balancing authorities or reserve sharing groups). Compliance is measured by analyzing data from events in a year through the annual Frequency Response Measure (FRM) (see [14] for details). Even though BAL-003-1 gives examples where load disconnection is discounted from IFRO requirements, in this work we propose that loads can be used to provide primary frequency response, and hence should be contributing towards the IFRO requirements. System-wide response of demand-side frequency responsive loads can be designed to complement generation response to achieve the required level of IFRO. That is, part of IFRO is provided by generation and the rest is provided by responsive loads. In the following, we propose a simple way of setting up parameters of controllable load to contribute to IFRO and FRO, in a similar manner as generators contribute.

Consider a power system structure where an interconnection is divided into areas 1 to N (or A_i with $i = 1, \dots, N$), containing loads and generation. In the case of the WECC system, the areas can correspond to Balancing Authorities, or Reserve Sharing Groups. Balancing Authorities group loads and generation in a geographic area, and Reserve Sharing Groups include several Balancing Authorities. The desired system-wide load response is defined as,

$$K_{sys} + IFRO_{sys-GEN} \geq IFRO_{sys} \quad (1)$$

where K_{sys} is the desired system-wide load control gain (MW/Hz), $IFRO_{sys-GEN}$ is the IFRO contribution from generators and $IFRO_{sys}$ is the system-wide IFRO response. Each area also contributes to the desired system response. The desired area-wide response is defined as,

$$K_{Ai} + FRO_{Ai-GEN} \geq FRO_{Ai} \quad (2)$$

where K_{Ai} is the desired area-wide load control gain (MW/Hz), FRO_{Ai-GEN} is the IFRO contribution from generators in the i -th area and FRO_{Ai} is the FRO for the i -th area. Also, the sum of the contributions from

each area should equal the expected system wide response $K_{sys} = \sum_{Ai=1}^{Ai=K} K_{Ai}$, and the sum of bus response

should equal the area response $K_{Ai} = \sum_{Busi=1}^{Busi=n} K_{Busi}$. One possible way of distributing the desired gains from the system level to the areas, and subsequently from the area level to the buses, is by making them proportional to the available controllable load as follows. The contribution of each area is made proportional to the controllable load of that particular area with respect to the total controllable load in the system using the factor $\frac{P_{cont_{Ai}}}{P_{cont_{Tot}}}$. Similarly, the contribution of each bus is made proportional to the

controllable load of that particular bus with respect to the total controllable load in the area using the factor $\frac{P_{cont_{Busi}}}{P_{cont_{Ai}}}$. The control gains for each area and at each bus can be defined as,

$$K_{Ai} = K_{sys} \frac{P_{cont_{Ai}}}{P_{cont_{Tot}}} \quad (3)$$

$$K_{Busi} = K_{Ai} \frac{P_{cont_{Busi}}}{P_{cont_{Ai}}} \quad (4)$$

2.3 Device-level control model

This work focuses mainly on thermostatically controlled loads (TCLs), in particular, water heaters (either electric resistance or heat pump water heaters), and HVAC systems. Such devices consume a non-zero power while in the ON state, and are assumed to be zero power in the OFF state. Hence, the total controllable power at one bus in the system is given by the sum of powers of all the ON TCLs that are directly connected to that particular bus. Each controllable load is expected to make an independent decision regarding its state of operation, such that the bus aggregated load response would match the desired power modulation. To accomplish this, a Markov-Chain model was adopted to capture the aggregated behavior of the controllable loads.

The simplest one is a 2-state model, representative of an electric resistance water heater, as depicted in Figure 2. The electric resistance water heater could be exclusively in one of two operating states, that is, either ON or OFF. This modelling algorithm ignores the natural transitions between states for individuals during a control period. For simplicity, it is assumed that switching between states occurs instantaneously. The state transition probabilities of ON (μ_{01}) or OFF (μ_{10}) are calculated so as to match the bus aggregated load power.

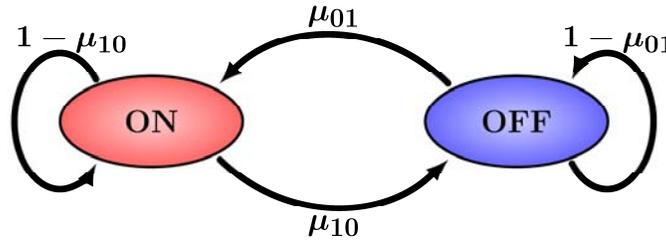


Figure 2. The 2-state Markov-Chain model of an electrical water heater

A heat pump water heater, HVAC, or other load with a motor is modeled using four states, as shown in Figure 3. In addition to the ON and OFF states similar to the electric resistance water heater model, two states (ON LOCKED and OFF LOCKED) are added. The lock periods, typically a few minutes, protect the compressor motor from wear and tear due to excessive cycling.

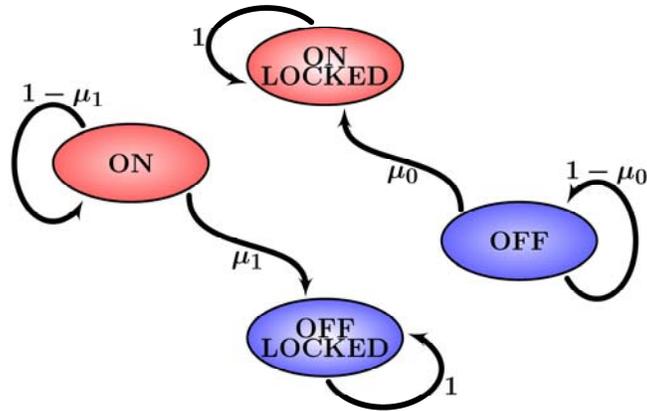


Figure 3. The 4-state Markov-Chain model of a heat pump water heater or HVAC system

3.0 Study of system-wide impacts of large-scale deployment of frequency-responsive load control

Large power system interconnection studies are typically performed by utility and research engineers with commercial software, like Siemens PTI PSS/E [16], GE PSLF [17], and more recently using the PowerWorld Simulator [15]. Utilities and interconnection coordinators, such as the WECC, update and manage the databases for their models. Most of the models are standard and are available in commercial software libraries. New proposed models and control strategies, like the hierarchical demand-side frequency control proposed in Section 2, should be incorporated in commercial tools as user-defined models (UDM), in order to be tested and validated. These new models remain as specific UDMs until the model becomes more widely used, after which the models can be incorporated into the standard libraries of the software packages. This section describes the implementation improvements made to the model used for the hierarchical demand-side frequency control on the WECC interconnection model. The following features are added to the previously developed user-defined model (UDM):

- The UDM was modified to interconnect with the transmission network through the WECC composite load model, which explicitly details the substation transformer and equivalent representation of distribution feeders, as well as the distinctive features of the individual load components which are represented by motor loads, power electronic loads, and static loads.
- To facilitate large scale deployment of the model, area-wide gain settings can now be input into the model. That is, rather than introducing a multitude of models with similar parameters, groups of distinctive models were created for each area in the transmission model, increasing the heterogeneity of the frequency responsive devices.
- The UDM models aggregated water heaters and HVAC systems that have been connected to the static load component of the WECC composite load model.

The model implementation is ensured to be correctly initialized to preserve the steady state solution returned by the power flow solver in PowerWorld by loading the steady state active and reactive power and voltage from the transmission simulator side. The improved UDM was used to test the hierarchical

gain design, described in Section 2.2, to meet NERC standards for frequency response obligation at both the area and interconnection levels.

3.1 Modeling hierarchical demand-side frequency control in PowerWorld

Dynamic models of large power system interconnections, such as the WECC, are composed of individual dynamic device models, such as synchronous generators with their controls (e.g., automatic voltage regulators and governors), coupled with a model of the transmission network. Load models include aggregated representation of many end-use loads and the distribution networks that service those loads. The representation of such models can either be static, such as the ZIP model (portions of constant impedance, constant current, and constant current) or dynamic, such as motor load models. Currently, new load models have been implemented in the WECC system model as a composite load model, as described in [18], [19] and [20].

The WECC composite load model includes various types of motors, electronic loads, static loads, and an equivalent of the distribution network and substation transformers, as discussed in [18], [19], and [20]. Previous work under this project [2] did not consider the newly developed WECC composite load model. The user-defined model (UDM) implemented in [2] was a static ZIP load model. Using the static ZIP load model could be an acceptable approximation to represent water heater loads. However, it is a crude approximation for HVAC loads, where motor dynamics need to be included. Furthermore, ZIP load models do not consider the effects of the distribution network between the high-voltage substation and the end-use load. For example, when resistive water heater loads at the end-use change, the reactive power consumption of the distribution network changes, causing a change in the reactive power load at the substation level.

The UDM for the hierarchical demand-side frequency control strategy proposed in Section 2 consists of the following elements: a) load dynamic behavior, which is captured by the different versions of Markov chains described in Section 2.3.2 of [2], representing dynamic behavior of populations of end-use water heaters and HVACs; b) a frequency controller that affects the water heaters and HVAC populations, as discussed in Section 2.4 of [2]; and c) an electric load model that is modified by the previous two elements. In FY15, (this report), the UDM of [2] was interfaced with the WECC composite load model, as mentioned before. Figure 4 shows a diagram of the UDM with the WECC composite load model in PowerWorld software. The user-defined model requires the current values of bus frequency (for control) and voltage for (load modeling) as input from the power network. Based on these frequency and voltage, the UDM models load dynamics representing the hierarchical demand-side frequency control and the natural load dependence on voltage. The updated values of loads are communicated to the rest of the power system model at each time step of simulation.

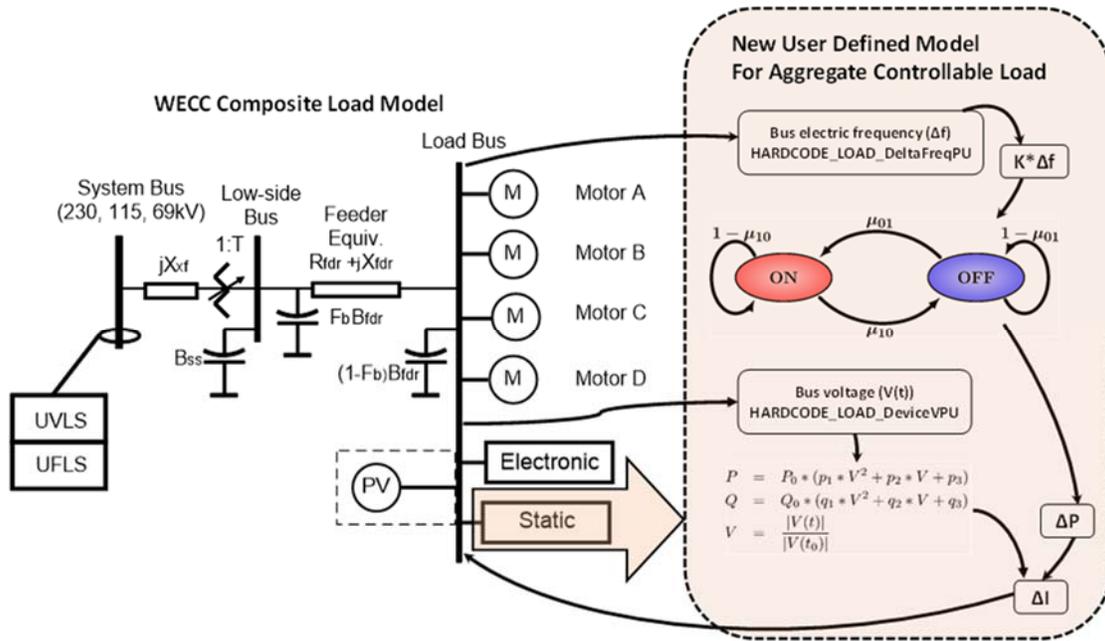


Figure 4 Diagram of user-defined model (UDM) of frequency-responsive load controller in PowerWorld and its interface with the new WECC composite load model [19]

The interface, initialization, and large-scale deployment of the UDM are discussed further here. As mentioned before, the load UDM communicates with the PowerWorld simulator to access voltage and frequency data. This is achieved with the following hard coded signals from PowerWorld: `HARDCODE_LOAD_DeltaFreqPU`, which represents the load bus frequency deviation from nominal (60 Hz), and it is used by the device model to initiate correct actions due to under-frequency events; and `HARDCODE_LOAD_DeviceVPU`, which represents the bus voltage magnitude in per unit (p.u.), and it is used to calculate the correct current injection for the current bus load. To ensure correct initialization of the load UDM, the bus steady state solution returned by the simulator power flow solver needs to be loaded as initial parameters into the load model, which is an improvement to the current model. Hence, the initial load is set based on the steady states of the bus active and reactive powers. Moreover, to capture dynamic load behavior, a simple ZIP load model was written inside the UDM that captures the voltage dependency of some loads, such as resistive water heaters.

For large-scale deployment of the UDM in a large complex system model like the one for the WECC, there is a challenge with creating a heterogeneous population of load models spread across the entire network. To resolve this challenge, groups of frequency-responsive UDMs were created for each system area. Thus, at most 21 different UDM types have been introduced into the system. As a future endeavor, creating an auxiliary file built by a scripting language is considered as one solution to increase model heterogeneity.

3.2 WECC test scenarios

The WECC 2015, high-load summer case was used to test the hierarchical demand-side frequency control connected to the WECC composite load model and to show contribution of the controller to meet NERC standards. This section provides a summary of the WECC system model and the scenarios of controllable loads in terms of amount and location.

The WECC system has 20,573 buses, 3,984 generators, 16,426 transmission lines, and 10,922 loads. As shown in Figure 5, the WECC system covers the western parts of the United States, Canada, and the northern portion of Baja California, Mexico. The 2015 high-load summer case has a total load of 159,104 MW. To study the frequency response, a contingency is simulated where large generating units are tripped in the southern part of the system. The WECC system consists of several balancing authorities, as shown in the right side of Figure 5. The balancing authorities are also grouped into reserve sharing groups (left side of the figure). Therefore, there are at least five possible levels of coordination in WECC system: interconnection level, reserve sharing group, balancing authorities, substations, and individual devices.

In this report, we illustrate the coordination, showing how controllable loads can contribute to BAL-003-1 requirements by supplementing the MW/Hz response provided by generators. Also, we look into the response of a particular area and see how the MW/Hz of response can be incremented, while simultaneously contributing to the BAL-003-1 requirements at the area level. The IFRO requirement for WECC is about 8400 MW per 1.0 Hz at settling frequency [3]. The WECC model utilized in this example exceeds this requirement, with a response of about 31,000 MW/Hz, as will be shown in Section 3.3.1. Instead of focusing on reaching the 8400 MW/Hz of the requirement, the study of this report is setup to show how a specific additional MW/Hz amount of response can be achieved with controllable loads across the WECC system. This example studies the system-wide response and the response of one area, part of California/Mexico Reserve Sharing Group. The hierarchical demand-side frequency responsive control strategy was set up so that load resources provide additional 7,955MW/Hz response the WECC interconnection model. It was assumed that given a certain IFRO and a known required generator contribution, additional 7,955 MW/Hz (5% of the total steady state load of the WECC system, 159,104 MW) was needed to achieve required system frequency response. This value represents the controllable load contribution towards the IFRO, complementing the response from generators.

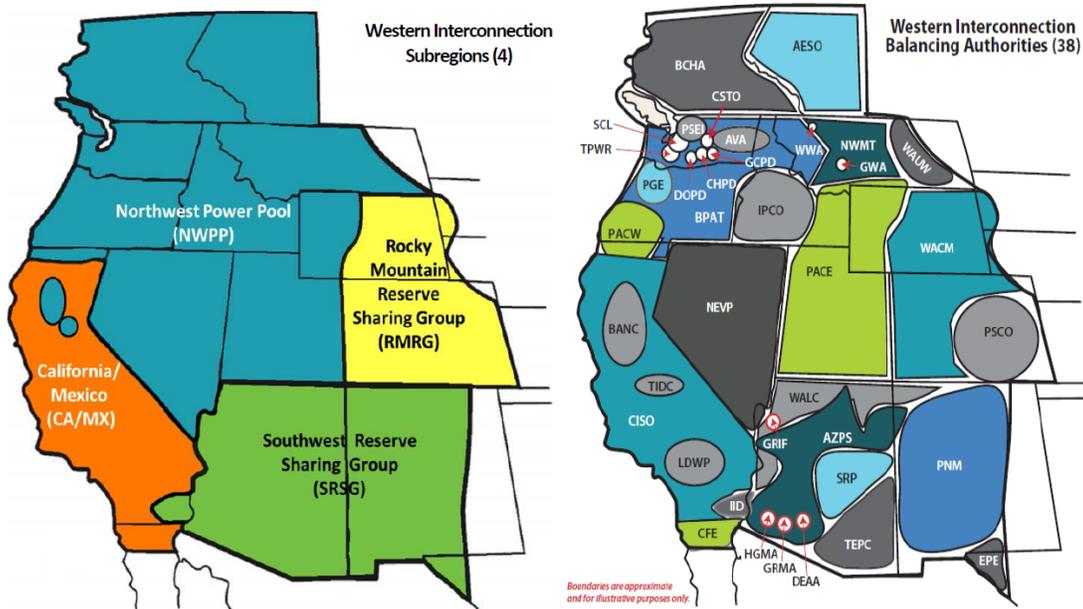


Figure 5. WECC balancing authorities (right) and sub regions for reserve sharing groups (left) [21]

3.3 Simulation results

Using the configuration of the WECC model discussed in Section 3.2, this section shows the results of simulation studies to demonstrate the efficacy of the hierarchical demand-side frequency controller in meeting the NERC BAL-003-1 requirements. Three test cases were considered. In Case 1 there was no controllable load. In Case 2, controllable load of 12,617 MW was uniformly distributed across a total of 5,723 buses, with a system-wide gain of 7,955 MW/Hz. The southern part of the system had 1,982 MW of controllable load with a gain of 1,250 MW/Hz. In Case 3, controllable loads were concentrated in the southwest portion of the system, where the total controllable load of 5,934 MW is distributed across a total of 1,813 buses, with a system-wide gain of 7,955 MW/Hz. In the same southern part of the system as Case 2, there was 1,982 MW of controllable load with a gain of 2,657 MW/Hz. Table 1 shows the settings used for controllable loads in the WECC system model for all 21 areas. These gains were calculated based on the procedure discussed in Section 2.2 based on the FRO requirements for the different areas. The 21 areas of the WECC system model in Table 1 (which are traditionally used in the WECC databases) are an aggregation of the 38 balancing authorities shown in Figure 5.

Table 1: Summary of controllable load settings for all 21 areas of WECC system model

Area	Total number of user-defined models (UDM) attached	Area total controllable loads (MW)	Gain for controllable loads (MW/Hz)	
			Case 2	Case 3
1	505	725.75	457.58	0
2	277	951.89	600.17	0
3	41	49.10	30.96	0
4	122	296.97	187.24	0
5	89	132.41	83.48	0
6	140	356.72	224.91	0
7	1,409	2,185.57	1,378.01	0
8	388	564.96	356.21	0
9	1,084	1,459.32	920.10	1,956.37
10	313	1,134.45	715.27	1,520.85
11	91	101.54	64.02	0
12	174	159.35	100.47	0
13	42	68.10	42.93	91.29
14	82	627.18	395.44	840.81
15	76	175.74	110.80	0
16	98	534.02	336.70	715.92
17	85	128.59	81.07	172.39
18	109	1,982.40	1,249.91	2,657.62
19	329	589.32	371.57	0
20	266	391.71	246.97	0
21	3	2.25	1.42	0
System-wide totals	5,723	12,617.31	7,955.25	7,955.25

3.3.1 Analysis of results for system-wide frequency response

Figure 6 shows the frequency at a representative bus where controllable loads are located. In both Case 2 and Case 3, the frequency deviation was arrested faster. The frequency deviation also recovered to a value closer to nominal, as compared to Case 1, when there were no frequency responsive loads. The frequency behavior in Case 2 and Case 3 is very similar due to the fact that the desired frequency response for the overall system, a load reduction of about 7,955MW for a 1Hz frequency drop, is the same. Cases 2 and 3 were set up with the same overall gain, but with different distributions of controllable loads.

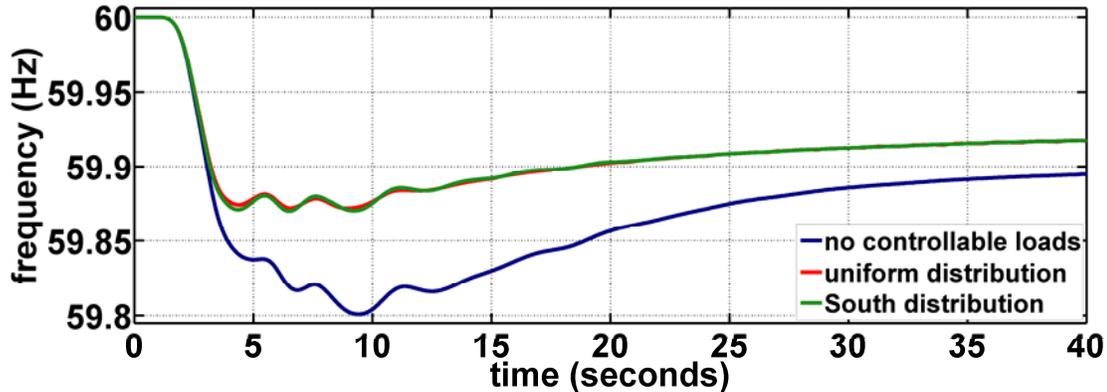


Figure 6. Frequency at a bus in the Northwest area for cases 1 (no controllable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)

The values in Figure 7 at 40s, when the frequency is in the steady state, show that the total system frequency response (the contribution from both generators and loads) is roughly 37,000 to 38,000 MW/Hz, while the case without controllable loads (contribution from generators only) had a little over 31,000 MW/Hz. Frequency response load increased the system IFRO by 7,000 MW/Hz and 6,000 MW/Hz for cases 2 and 3, respectively. The full 7,955MW/Hz is not achieved because of the voltage dependence of the loads. As a result of load disconnection in response to the frequency deviation, voltages increased as compared with the case with no control. This voltage increase counteracts some of the load reduction achieved by the controllable load.

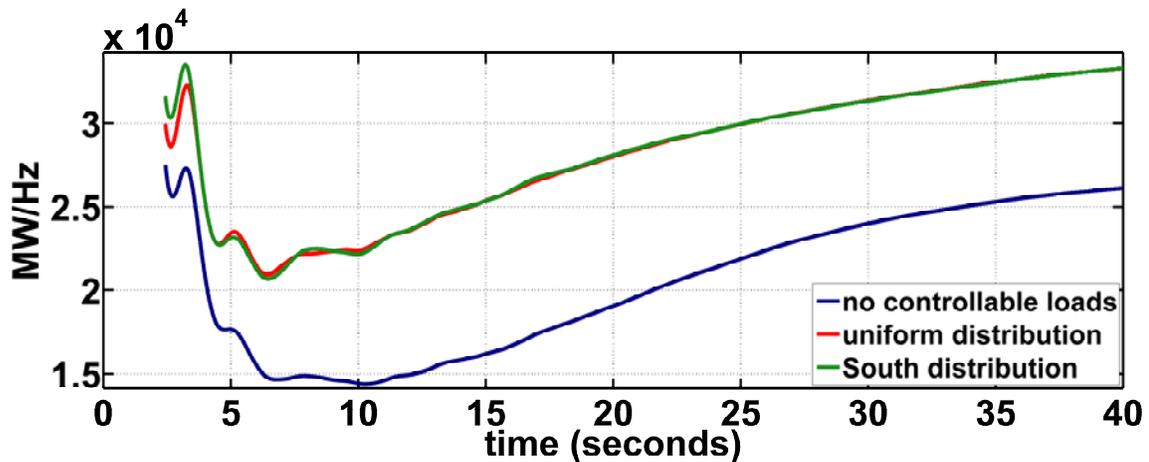


Figure 7. System overall frequency response for cases 1 (no controllable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)

Additionally, using controllable loads as resources to control the frequency alleviated the stress on the important tie lines in the WECC system interconnection. As seen in Figure 8, the tie line carrying power from north to south increases after the contingency (generation trip in the south). These results are similar to the cases when there is no controllable load and when the controllable load is spread uniformly across the system. However, when the controllable load is concentrated in the south, closer to the area where the contingency occurred, the tie line power is substantially reduced. The contribution of generation, loads, and losses to the overall system response of Figure 7 can be discussed further.

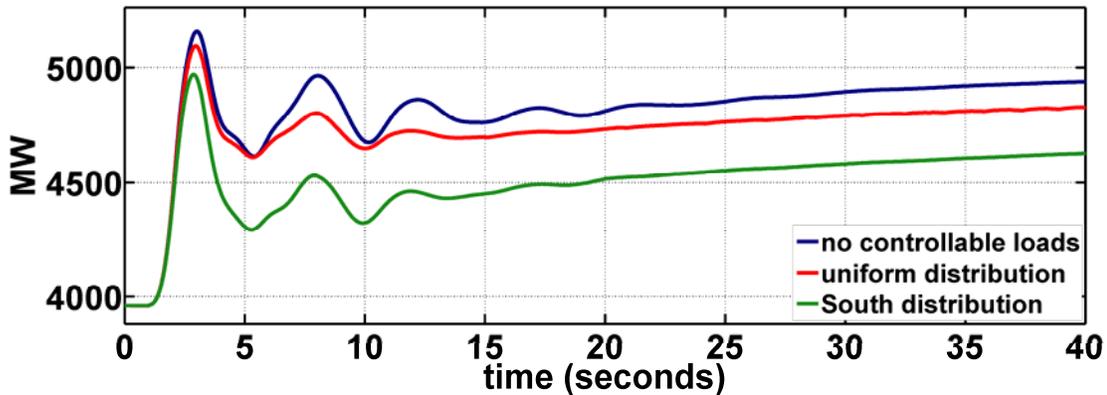


Figure 8. Tie line power flow, north to south, for cases 1 (no controllable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)

Figure 9 shows that the contribution of generators is very similar in the cases with controllable load (Case 2 and 3), and it peaks higher in the case of no controllable load. The values of generation contributions at 40s represent 95% of the overall frequency response for the case with no controllable load, and about 70% of the overall frequency response in Case 2 and Case 3.

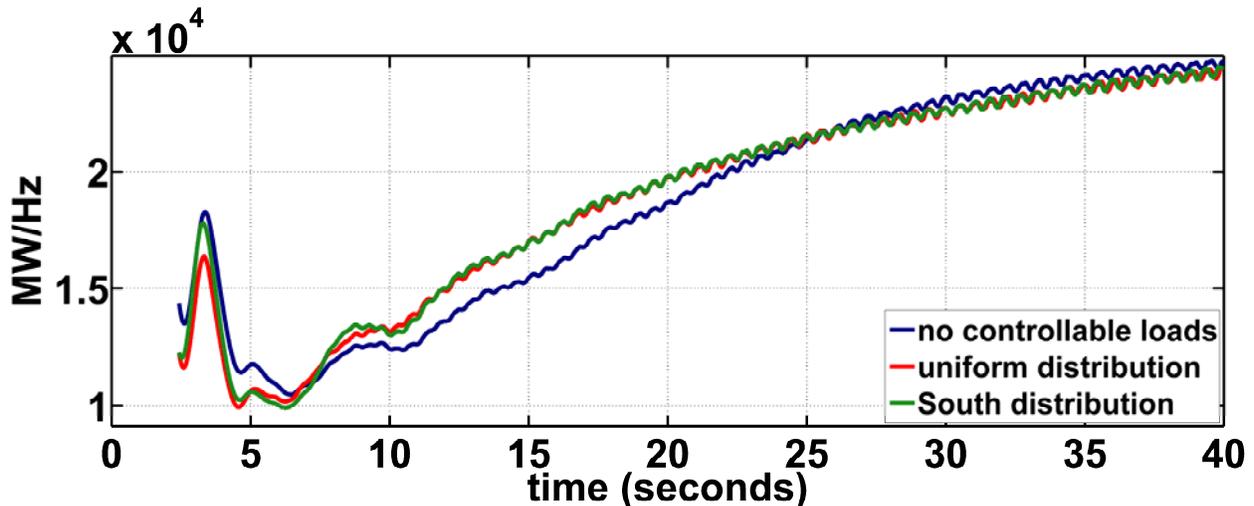


Figure 9. Contributions of generators to overall frequency response for Cases 1 (no controllable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)

Figure 10 shows the increased load contributions for Cases 2 and 3, with respect to the natural response of the loads in the case with no controllable loads. The values of load contributions at 40s represent 5% of the overall frequency response for the case with no controllable load (natural response), but load contributions are 41% and 37% of the overall frequency response for the Cases 2 and 3, respectively.

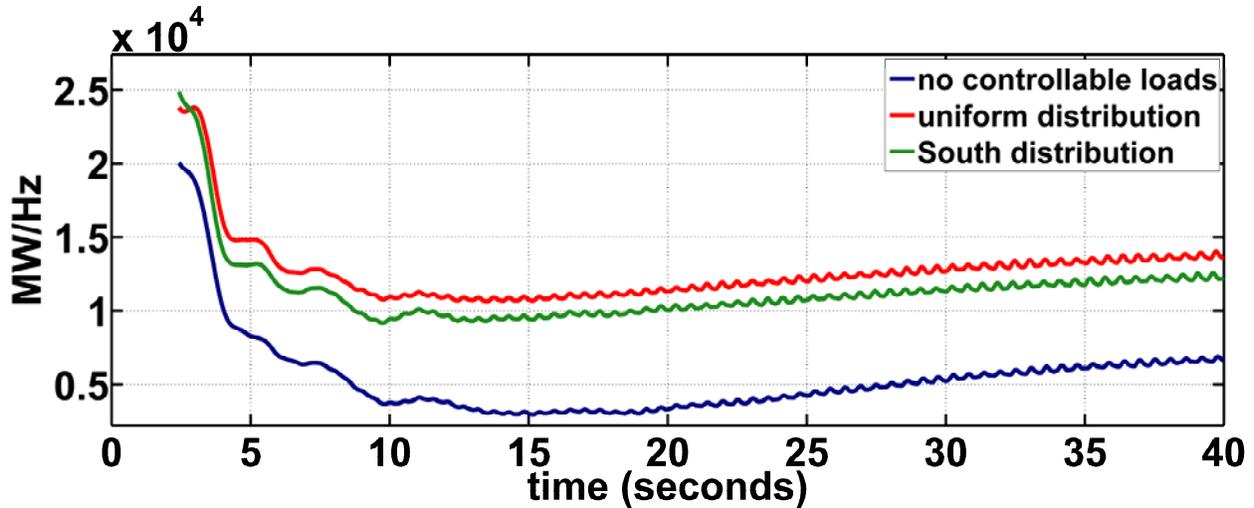


Figure 10. Contributions of loads to overall frequency response for Cases 1 (no controllable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)

Finally, Figure 11 shows the contributions of the losses. The losses contribute negatively to the response of the particular contingency under study (generation loss in the south). This is because the north-south power flow increases after the generation is lost in the south, increasing losses as extra power is delivered from the northern portion of the system. It can be seen in Figure 11 that the negative contribution of the losses is reduced for the case with controllable load concentrated in the south (Case 3). The values of losses contributions at 40s represent -20% of the overall frequency response for the case with no controllable load. Loss contributions are -14% and -10% of the overall frequency response in the Cases 2 and 3, respectively. It can be observed that the effect of the load contribution compensates the effect of losses contribution to obtain the same overall system response of Figure 7.

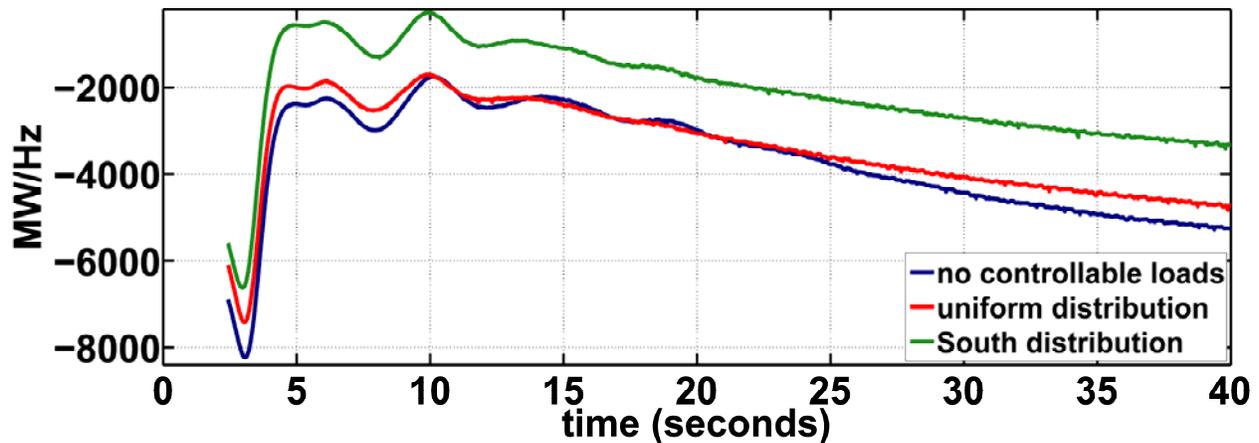


Figure 11. Contributions of losses to overall frequency response for Cases 1 (no controllable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)

3.3.2 Analysis of results for area-wide frequency response

From an area perspective, the contribution towards a system-wide value of 7,955 MW/Hz IFRO is dependent on the distribution of the frequency responsive loads. Concentrating the frequency responsive

loads in particular areas in the South resulted in a higher control gain for those areas to be able to achieve more frequency response, similar to the case with uniform distribution (Case 2). The controllable load response contributes towards the FRO area requirement (as detailed in Table 1), while adequately meeting the IFRO. Figure 12 shows that for the case without controllable loads (Case 1), a steady state (value at 40s) of 2,000 MW/Hz load contribution to the FRO was achieved (this response comes mostly from generation). When controllable loads were distributed uniformly system-wide (Case 2), an additional 1,000 MW/Hz approximately was added (total area response of 3,000 MW/Hz). Concentrating the controllable loads in the southwest portion of the system (Case 3) led to an additional 2,000 MW/Hz. This final scenario resulted in a total area response of 4,000 MW/Hz (with an additional 2,000 MW/Hz over the case 1, with no load control).

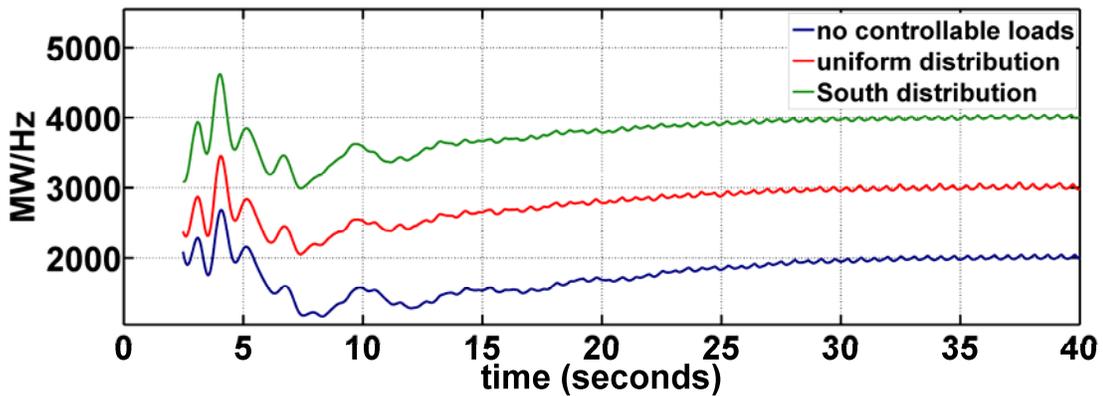


Figure 12. Area overall frequency response of an area in the south for Cases 1 (no controllable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)

It is also worth discussing the changes in area generation and loads due to the different distribution of controllable loads in the South. Figure 13, shows that the area generation is reduced when the load control is added (compare response from Case 1 to responses from Case 2 and 3). It is interesting to observe that the generation response is the same for the two cases with controllable load (Cases 2 and 3). In other words, only the overall system response from the loads affects the generation burden on frequency response. However, it can be seen in Figure 14 that the area load does decrease for both Case 2 and Case 3, as expected. The area load decrease for Case 3 (controllable load concentrated in the south) is twice as much as compared to the reduction for Case 2 (load gain uniformly distributed across the WECC). This is due to concentrating the controllable load in the same part of the simulated WECC model where the contingency occurs, that is, the southern portion, rather than being uniformly distributed all over the system.

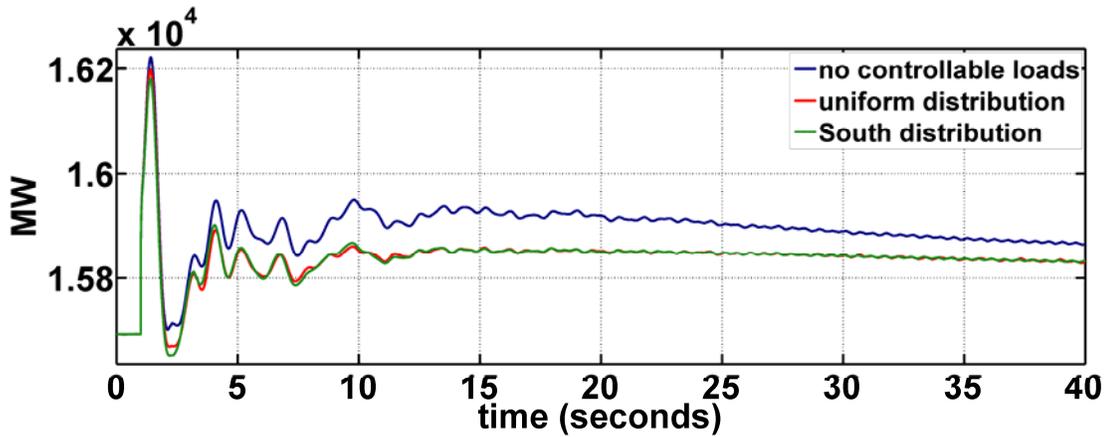


Figure 13. Total area generation of an area in the south for cases 1 (no controllable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)

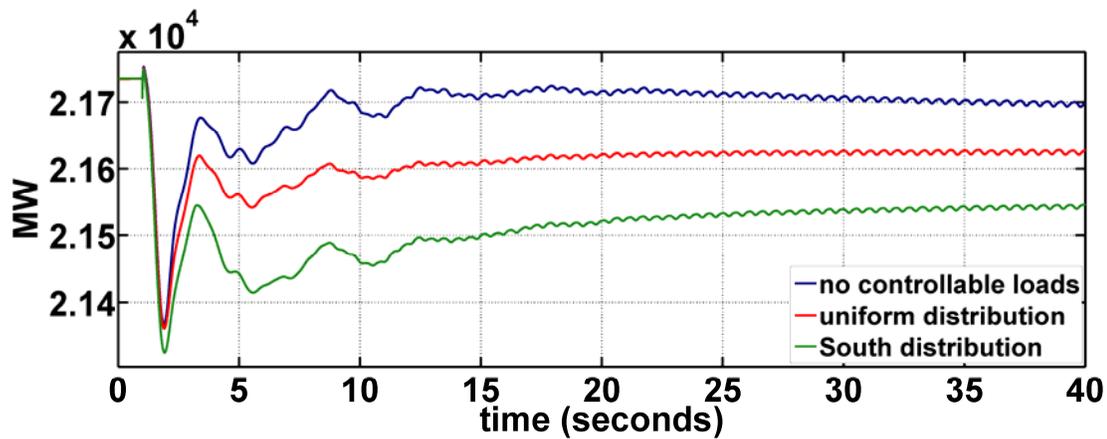


Figure 14. Total area load of an area in the south for cases 1 (no controllable load), 2 (uniform distribution of controllable load) and 3 (controllable load concentrated in the south)

The results in this section illustrated how the autonomous response from water heaters can be coordinated from a supervisory level to obtain desired response in terms of MW/Hz. This response is obtainable at both the area and system levels, using the simple aggregate load models introduced in Section 2.3. In Section 4.0 hardware testing is performed to validate and calibrate the water heater models. Improvements will be made to individual models, as needed, in order to better capture the observed behavior of the physical hardware.

4.0 Hardware testing of proposed frequency-based load control strategies

In this section, the results of the hardware testing and model validation task are presented. The main objective of this test was to verify the applicability of the control logic over a wide range of operating conditions while accounting for unpredictable end-use device behavior and physical device constraints. Testing was carried out on residential-scale water heaters, appliances that have been identified as candidates for frequency control. This was done to verify that the expected response to the control signal is carried out and to validate and improve the appliance models to be used in large-scale population studies. Individual behavior was first verified for individual appliances. Load composition studies were leveraged to extrapolate observed behavior to scale to a large population. Changes were made to aggregate models in order to better capture the observed behavior of the physical hardware and population characteristics.

Existing laboratory resources at PNNL were leveraged for this task. Testing was carried out on two large tank water heaters, an electric resistance water heater (ERWH) and a heat pump water heater (HPWH). The performance of the HPWH was compared to the ERWH to design performance requirements for devices at the individual level in providing primary control response desirable at the aggregate level. In addition, the total capacity and performance of primary frequency control of the aggregate population of water heaters compared to the system scale was determined for both types of water heaters.

4.1 Hardware setup used for experiments

Thermostatic loads are able to provide grid support services with negligible impact to the customer because their thermal inertia can sustain the system for short periods of time when power is removed without affecting their provision of the intended service. Residential loads such as water heaters, air conditioners, electric dryers, and refrigerators all fall into this category. Water heaters are ideal because they have significant stored thermal energy and constitute 18% of the total residential load, a sizeable fraction [22]. Each unit has a relatively high power consumption (on the order of kW when heating is engaged, one of the largest of typical residential loads) and there is a large installed base.

Historically, water heaters have used a resistive heating element and relied on the action of a purely mechanical thermostat to change state, with no solid state logic. New models of water heaters that use a heat pump are increasingly becoming available, and their penetration is likely to increase markedly in coming years with the adoption of a new federal standard for large-tank water heaters, 75 FR 20112 [23].

This study leveraged an existing laboratory set-up that was created for the purpose of studying and comparing the performance of traditional demand response at the timescale of hours to more rapid modes at the timescale of seconds [24]. Each experiment in this study included testing of an 80-gallon class ERWH and HPWH, with the same scheduled draw pattern and water temperature set point. The scheduled draw pattern was varied to capture the impacts of typical draw patterns on water heater performance. The two water heaters were installed in a laboratory space, as shown in Figure 15.

Monitoring and control equipment was rack-mounted. The ambient temperature, relative humidity, and inlet water temperature were controlled by the building's environmental control systems.

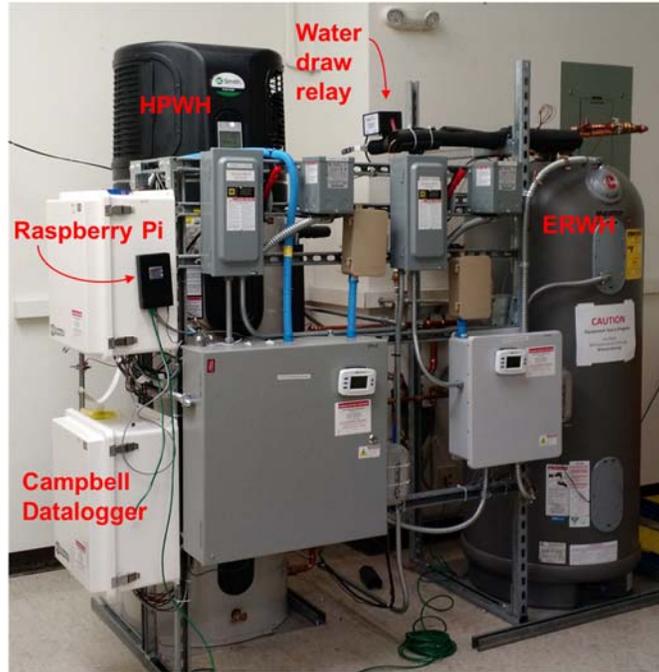


Figure 15. Hardware configuration in laboratory.

The HPWH was the 80-gallon A.O. Smith Voltex® Hybrid Electric HPWH, model# PHPT-80, the most common large-tank (>55 gallon) HPWH for residential use. The ERWH was the 85-gallon Rheem Marathon ERWH, model# MR85245, commonly used in existing demand response programs. The water heaters and primary sensors were plumbed in accordance with ASHRAE Standard 118.2-2006, “Method of Testing for Rating Residential Water Heaters” [25]. A single supply of cold water was split with a tee connection to supply both water heaters. Based on a field study of hot water use data [26], both water heaters were configured with 125°F temperature set point and 1.66gpm flow rate. Measurements were taken of the parameters listed in Table II and data recorded at 100ms intervals. The monitoring system included two Campbell Scientific model CR1000 data acquisition systems. Prior work [24] did not implement controls that required updated operational parameters. For this study, monitoring of power and water temperature was achieved by network-connecting the Campbell data acquisition systems to the internal PNNL network, which allowed the measurements to be queried by control software at 100ms intervals. More detail about the data acquisition system design and the temperature sensor physical locations can be found in [24].

Table II. Monitored Parameters

ERWH	HPWH
Power total (W)	Power total (W)
Power resistance element 1 (W)	Power compressor (W)
Power resistance element 2 (W)	Power fan (W)
Flow (gpm)	Power resistance element 1 (W)
Temperature at 6 levels (°C)	Power resistance element 2 (W)
T_{ave} (°C)	Flow (gpm)
	Temperature at 6 levels (°C)
	T_{ave} (°C)

The control and experiment data flow is shown in Figure 16. A control agent was developed in VOLTTRON, an agent-based communication and control platform written in Python [27]. The network-connected Campbell Scientific CR1000 data acquisition system digitized the analog information, such as power and temperature on each water heater. This data was sent to a VOLTTRON agent that translated the CR-1000 data into the VOLTTRON data format and published on the VOLTTRON message bus. The VOLTTRON control agent subscribed to the water heater data, processed the data and responded to frequency read from a file using its control algorithm. The control agent then published the HPWH, ERWH and valve control commands as topics on the message bus. The hardware agent subscribed to the valve control, ERWH control, and HPWH control topics on the message bus; implemented temperature set point and configurable dead band control; and implemented this control through the hardware driver. The hardware driver was written in C, which included low-level signals needed to open and close heating element and flow control valve relays. The controllers were component level hardware relays, optocouplers, and other devices needed to implement the control logic on the physical hardware devices. The device driver, control agent, and flow control agent were all loaded onto a networked Raspberry Pi.

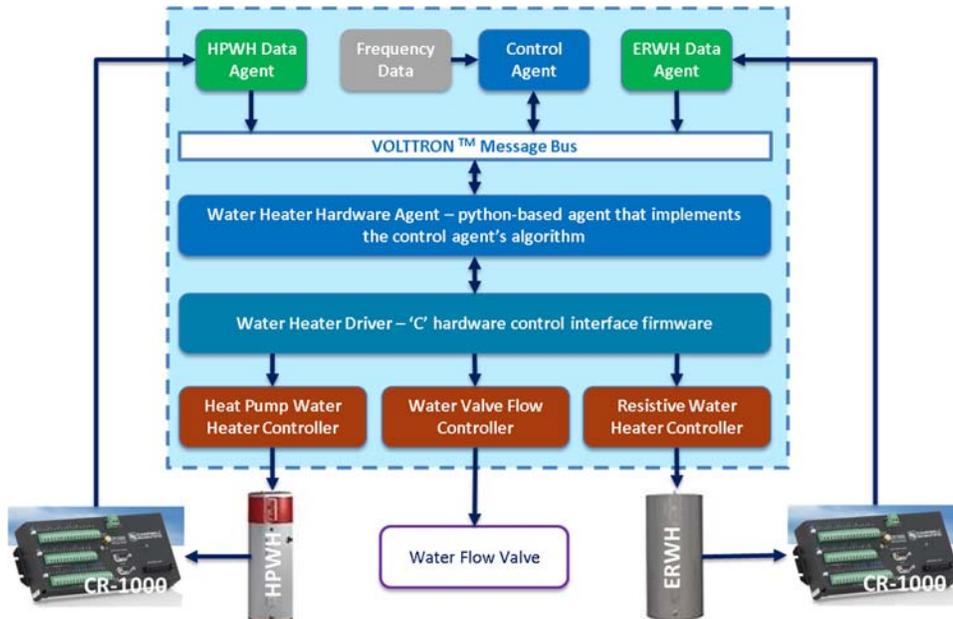


Figure 16. Control and monitoring system implementation

Separately, a schedule was implemented on the flow control relay, in order to represent typical water usage. An illustrative water draw profile over a single 24-hour period and resulting power use for the ERWH (top panel) and HPWH (bottom panel) is shown in Figure 17. This baseline water draw (identical for the two water heaters) and power consumption profile does not include the response to frequency events, only the changes of state that are necessary to maintain water temperature at the setpoint in response to water draw and thermal loss through the envelope. This baseline data was used to evaluate the probability that a water heater would be in a particular state in response to normal use.

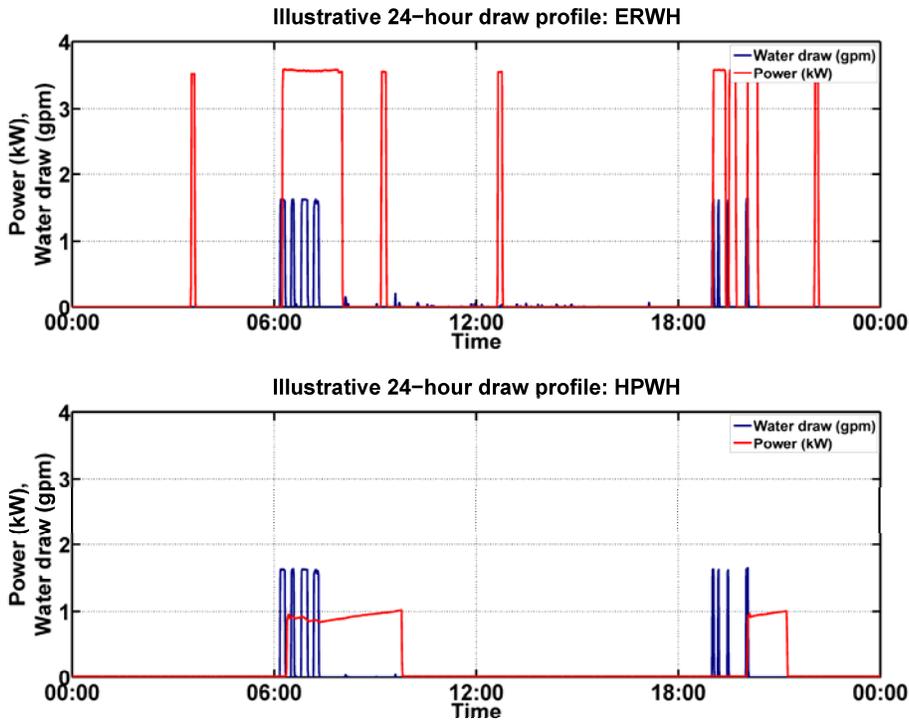


Figure 17. Water draw and power use over an illustrative 24-hour period.

In an ERWH, hot water is drawn from the top of the tank and cold make-up water is supplied near the bottom of the tank, and a stratification layer is maintained between colder water at the bottom of the tank and warmer at the top. There are two resistive heating elements, one located near the bottom of the tank and one near the top or about two-thirds of the way up, each responding to a thermostat located near the element. Under light to normal draw conditions, the lower heating element alone provides the heating. The upper element is used when the stratification layer moves past the position of the lower heating element. Only one 4kW element is powered at a time. The design allows for rapid re-heating of small amounts of water above the top element so that water is available at the temperature setpoint as quickly as possible. As can be seen in the top panel of Figure 17, the resistive heating element turns on shortly after a water draw of any duration begins, and may also engage if the temperature in the tank has drifted below the level of the lower deadband due to heat loss to the environment (as can be seen at 4:00 and 13:00). The thermostats in ERWH are entirely mechanical, with no solid state logic.

In a HPWH, there is both a heat pump heating mechanism and a resistive element, and multiple modes of operation. The only mode considered in this study is efficiency mode, in which only the 1kW

heat pump is used to heat the water, providing the highest efficiency and lowest cost of operation. The electric resistance element is used if the ambient temperature are such that the heat pump cannot be safely used (outside of 45°F – 109°F) or if the tank temperature drops below 59°F, and those conditions were not experienced in the indoor, temperature-controlled environment with the water draw profiles used in the study. The HPWH relies on two digital thermostats, one located near the top of the tank and one located near the bottom of the tank. When the temperature of the tank, as determined by the weighted average of the two temperatures, drops outside of the deadband about the setpoint, the heat pump is activated. Because the heat pump activates based on a measurement of the average temperature, rather than temperature near the colder make-up water, the heat pump may not become active until after significant water draw has occurred, as can be seen (at 20:00) in Figure 17. When the heat pump does engage, it is likely to remain in the ON state for a longer period of time than an ERWH.

4.2 Hardware testing and validation

This section describes studies that were carried out on the hardware water heaters and how results were combined with load composition studies to validate the models used in simulations.

4.2.1 Power Cycle Testing Results

The aggregate models used in simulations described in Sections 2 and 3 assume extremely simple behavior from the individual loads. Appliances are assumed to have two states, to draw a constant power in the ON state and to draw no power in the OFF state, with the transition between states occurring instantaneously, or within one simulation timestep of 8ms. A simple alternating change of state command at regular intervals was implemented to examine any non-ideal behavior that would contradict those assumptions. As shown in Figure 18, the observed ERWH power draw during a change of state from the OFF state to the ON and then back OFF matches the very simple model's behavior quite closely, to the limit of the resolution of the measurements. The observed power draw changed from 0W to the steady state of 3.89kW within a single 100ms measurement interval. The power draw may change on the order of a few percent over the course of several minutes as the ambient conditions and tank temperature change (as can be seen in Figure 17), but over the 5-minute timespan in the ON state, the power draw varied less than 0.25%. The transition from the ON state back to the OFF state was also completed within a single 100ms measurement interval. The HPWH performed similarly. Over the course of the few seconds relevant for the primary frequency control response, the behavior is approximated by the very simple device model with a high degree of fidelity.

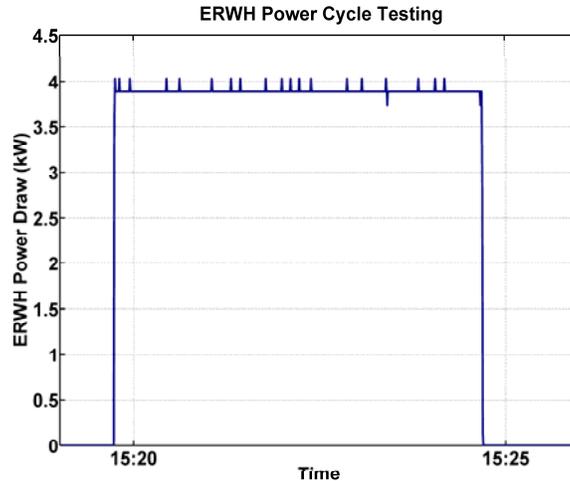


Figure 18. Power draw of ERWH through transitions of state.

The HPWH performed similarly, with the caveat that because the heat pump compressor is a motor, there is a lockout period: after a change of state, an additional change of state is not permitted within the lockout window. Lockout periods protect the motor by preventing excessive cycling. For this particular model, the lockout period is 7 minutes. This additional mode means that rather than the 2-state model which successfully describes the ERWH, a 4-state model is required to describe a HPWH. Both the 2-state model and 4-state model were introduced in Section 2.3. A 4-state load model was implemented in studies described in [2]. That load model was updated further and used to represent HPWHs, and a 2-state device model was updated and used to represent ERWHs. Similar motor-protective lockout periods can be seen in several other residential loads that have been identified as candidates for primary frequency control, including air conditioners, heat pump space conditioners, refrigerators, and pool pumps.

An additional important distinction is that, because the ERWH is a purely mechanical device, the change of state can be implemented in a very straightforward manner by interrupting power to the entire device. In contrast, the HPWH relies on digital thermometers and has digital control logic. If power is interrupted to the entire device, then a reboot process of two minutes must be undergone before resuming normal operations. For this reason, successful primary frequency control implementation using a HPWH similar to the one used in these studies must interface with the existing controls on the device at a lower level.

4.2.2 Water Heater Load Patterns

Load composition and characteristics has been a subject of study for decades, particularly in the WECC. One commonly used approach is to examine end-use survey data, such the End-use Load and Customer Assessment Program (ELCAP) conducted by BPA between 1984 and 1994 in the Pacific Northwest. Survey results can be coupled with laboratory component testing to capture detailed end-use device behavior. More recently, WECC's Load Modeling Task Force (LMTF) has developed composite load models and used recorded load data to validate the models at a system level [14]. Load model fraction is affected by ambient conditions (temperature, humidity, insolation, etc.), feeder characteristics, thermal characteristics of buildings, as well as human activities that vary by time of day, day or week, and time of year.

Typical water draw and water heater power demand profiles have been developed by a variety of studies. In 2012, the BPA carried out a Residential Building Stock Assessment [26], which collected detailed whole-house metering data for residences throughout the BPA service territory and used the data to develop load shapes for residential loads and assess the major determinants of residential energy use. As can be seen in Figure 19, the overall energy use varies +/-20% seasonally. This seasonal variation is second only to space heating for seasonal energy use changes, and indicates the necessity of a mechanism for updating devices' primary frequency control settings in order to achieve the desired aggregate response.

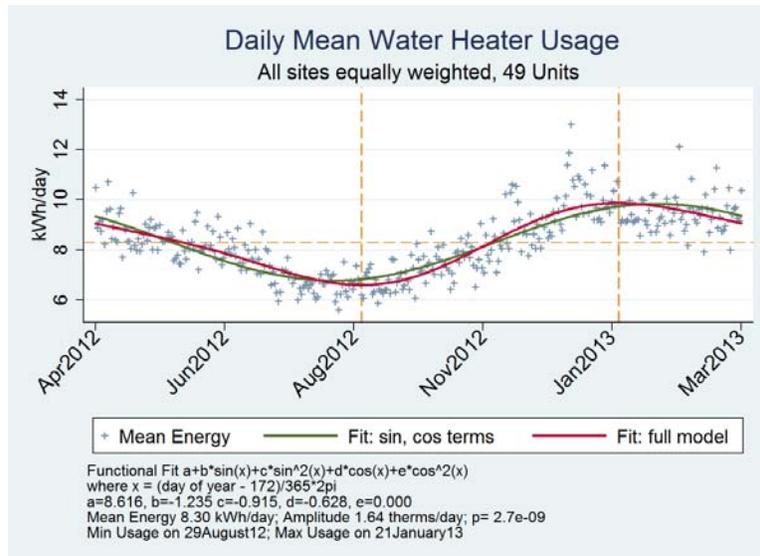


Figure 19. Daily Mean Water Heater Usage in the Pacific Northwest [28].

Water heater usage also varies over the course of the day, as can be seen in Figure 20. Individual sites (light gray lines) show a great deal of variation in hot water use patterns. Average demand (thick red line) is largest in the morning, as showers are typically the biggest use of hot water. Weekdays show a particularly pronounced morning peak, and weekends show a more temporally diffuse morning peak. A double-humped average load pattern is typical, with load peaking in the mornings and evenings. The RBSA data indicate that more clothes washing is done on the weekends, which results in an elevated but not strongly peaked contribution to the total hot water demand. Almost all residences surveyed used very little hot water between the hours of midnight and 5:00.

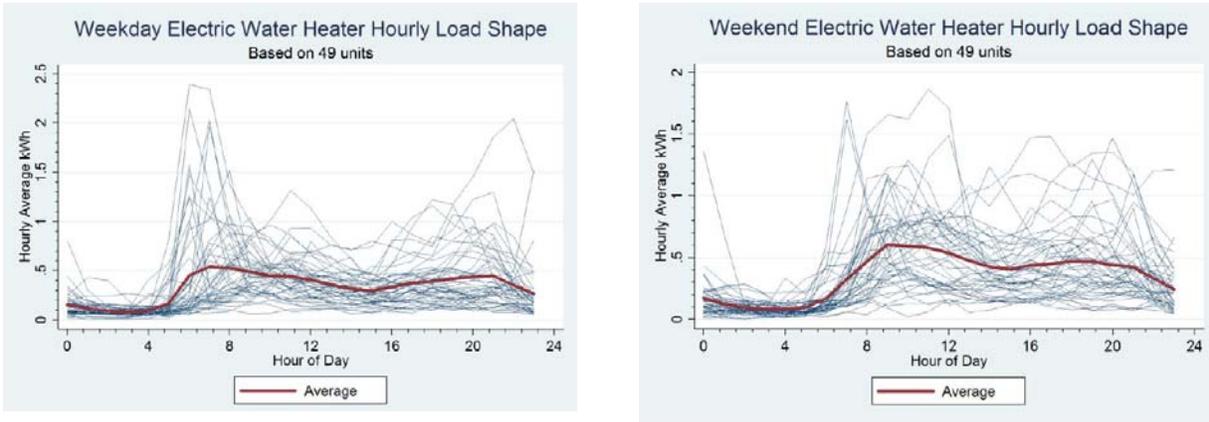


Figure 20. Water heater hourly load shapes in the Pacific Northwest [28].

The RBSA data develops a typical water draw profile, and load data can be used to infer a typical power demand from an ERWH. HPWH are newly available for residential settings and make up a very small fraction of the deployed total, so HPWH power demand patterns are not reflected in RBSA data. The water draw patterns, applied in a laboratory setting, as shown in Figure 17, can be used to develop a demand profile for a HPWH. A finding of note is that while both ERWH and HPWH show decreased water draws between the hours of 12:00 and 5:00, due to the averaging of low and high temperature measurements, the HPWH is significantly less likely than the ERWH to actively heat during those hours. Water heaters participating in primary frequency control would be able to contribute far less to the necessary response during nighttime hours than during daylight hours, with that trend exacerbating as HPWHs become more widespread. This observation indicates the importance of diversity in the loads that contribute to primary frequency control.

Following [24], in addition to the data from the RBSA, water draw profiles from the Building American House Simulation Protocols [29] and the Canadian Standards Association standard profiles [30] were also consulted. Based on this set of profiles, the average amount of time that the water heaters are found in each state was determined. For the ERWH, the heating element was in the ON state on average for 17% of a typical day, with a demand of 4kW when engaged. For the HPWH, the compressor was on average active for 22% of the day, with a demand of 1kW when engaged. In addition, the HPWH transitioned state on average 11 times per day; therefore 3% of the day was spent in the ON-LOCK state and 3% of the day was spent in the OFF-LOCK state.

4.2.3 Aggregate Model Validation

In order to assess how well the aggregate model used in the simulations discussed in Section 2 and 3 represented the behavior of physical devices under typical loading conditions, a multi-hour water draw pattern was set up, with the simulated frequency data read by the VOLTTRON control agent at a 10-minute loop. The time-series frequency data was recorded from a PowerWorld simulation of a disturbance in the WECC, as described in Section 3.2, with a simulated contingency where large generating units are tripped in the southern part of the system. The hierarchical demand-side frequency responsive control strategy was set up so that uniformly distributed load resources provide 7,955MW/Hz response in the WECC interconnection model. The frequency responsive load was simulated using either

the 2-state model UDM (for the ERWH) or the 4-state UDM (for the HPWH), and the UDM aggregated load with initialized with the populations in the states as described in Section 4.2.2 above.

The typical water draw profile caused changes of state, as indicated by the requirement to maintain the temperature at the setpoint. The frequency was read from the flat file by the VOLTTRON control agent, which treated it as measured frequency and implemented the control logic of the simulated loads. The results from 288 10-minute intervals were stacked in order to create an aggregate response from an effective population. A 10-minute loop was chosen in order to allow the HPWH to complete the 7-minute lock period that follows an induced change of state and return to the state that it would have been in due to water temperature alone. The results of these studies can be seen in Figure 21, for the 2-state ERWH, and in Figure 22, for the 4-state HPWH. The fraction of the population found in each state is plotted as a function of time. As in the simulations discussed in Sections 2 and 3, an event occurred 1 second from the start of the time interval, causing the frequency to deviate from nominal. Mismatches between the simulated and observed populations can be ascribed to two main factors: the difficulty of exactly matching the initial populations of each state between simulation and experiment, and the randomness of the Markov chain-based control logic, as implemented by a small number of devices. It can be seen clearly in both figures that the response of the hardware devices tracks the simulated aggregate model, with transitions between states occurring in response to changes in frequency.

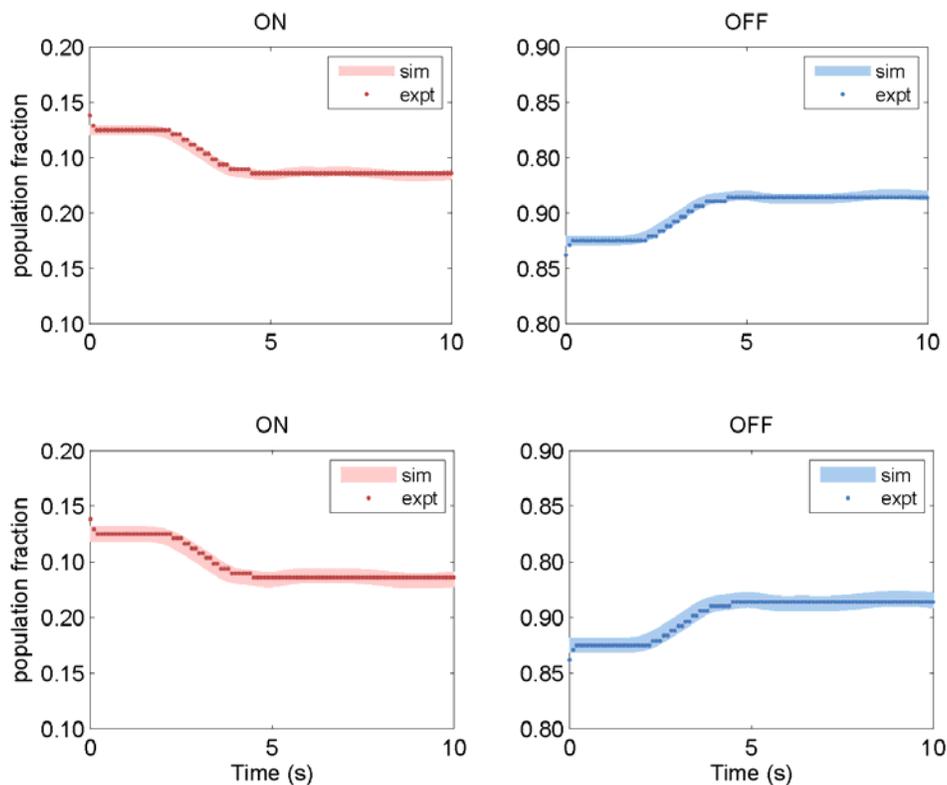


Figure 21. Simulated and observed state populations for ERWH

The use of the 4-state model rather than the 2-state model to represent the HPWH is necessary to capture the effect of the lock-state. Transitions between states occur only an average of 11 times per 24-hour period, and in this example, 22% of the population is in the ON state and 3% in the LOCK-ON state. If the populations in the two states were to be lumped together and considered as ON in calculating the

individual device settings, then the aggregate response would be 88% of the anticipated MW/Hz. However, the 7-minute lock time observed for the HPWH model used for testing is unusually large; many motors have protective lock times similar to 2 minutes.

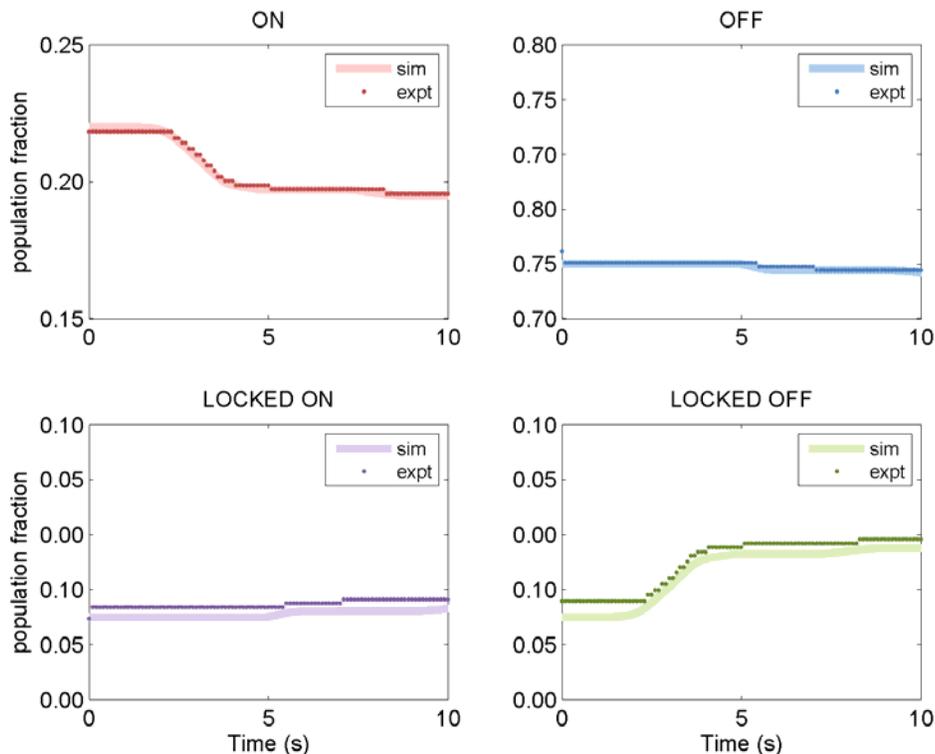


Figure 22. Simulated and observed state populations for HPWH

4.3 Population studies based on hardware testing

The results of the model validation studies discussed in Section 4.2 were used to update the supervisory layer calculations and aggregate load model introduced in Section 2.2 and 2.3 with detailed device characteristics. Realistic initial state populations were implemented, and the supervisory control layer was altered such that an appropriate gain could be calculated in order to attain a desired MW/Hz response given a population of residential water heaters of a particular type. The effect of varying penetration levels of frequency responsive load was studied using the same simulated scenario introduced in Section 3.2. The WECC high summer case was simulated in PowerWorld, with a contingency due to large generating units tripped in the southern part of the system.

Figure 23 shows the frequency as measured at a bus in the northern part of the WECC for varying penetration levels of primary frequency control-enabled ERWH. The base case has no controllable loads, relying entirely on governor controls in centralized generators. When only 10,000 or 100,000 ERWH participate, the resource is not sufficient to achieve a result significantly different from the base case. In that case, the population of ERWH in the ON state is quickly depleted, and the desired MW/Hz response is not achieved. A population of 1,000,000 ERWH shows an improved response, but again, the population of loads in the ON state is depleted before the target MW/Hz is attained. The 10,000,000 ERWH case shows the full desired response and since the resource is more than adequate, if the penetration of frequency-responsive ERWH is increased further, the frequency response is the same because the

supervisory control layer reduces the gains to provide the desired MW/Hz response by changing the state of a smaller fraction of the total population of participating ERWH. There are roughly 32,000,000 households in the WECC, and the pool of water heating devices that could participate in frequency control is not very much smaller. In that context, 5,000,000 is a high penetration level but not unattainable in the WECC for an individual end-use device.

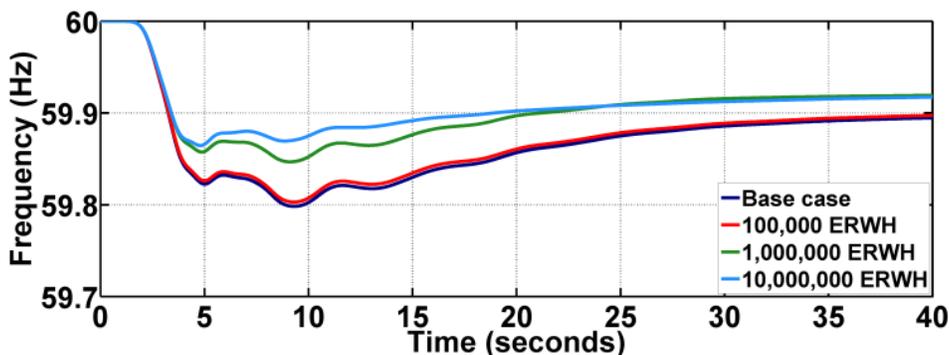


Figure 23. System frequency for varying participation levels of ERWH.

The performance of populations of HPWH in the same simulated scenario is shown in Figure 24. The base case has no controllable loads, and the frequency is not significantly different from the base case when there are only 100,000 HPWH participating in primary frequency control. Some improvement is seen at penetration levels of 1,000,000 HPWH and then the resource saturates (for this disturbance event) at less than 10,000,000 HPWH.

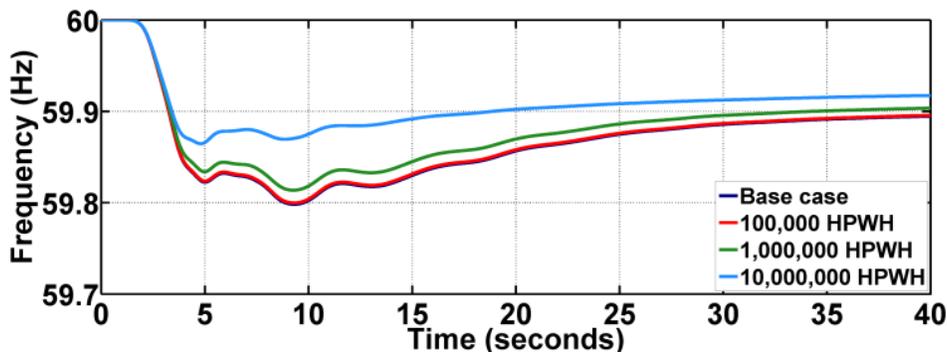


Figure 24. System frequency for varying participation levels of HPWH.

In order to quantify the effect of the controllable loads on the system frequency, the A-C performance metric is defined where the nominal frequency is referred to as ‘A’ and the nadir as ‘C’, and their difference as the A-C metric. As can be seen from Figure 25, when the population of frequency-responsive loads is lower than about 100,000, the impact of controllable load on the frequency is negligible. The controllable loads do not cause any negative impacts, but do not appreciably improve the response.

As the population grows to several million, it is increasingly able to make a difference in the observed frequency response, but the response may be difficult to predict since the resource will be prone to being depleted. While the HPWH are more likely to be found in the ON state than the ERWH (22% vs 17%), a HPWH in the ON state draws only an average of 1kW, in contrast to the 4kW of an ERWH. Therefore, a

larger population of HPWH than ERWH is required to provide the same MW/Hz response. The resource saturates at less than 5,000,000 participating ERWH for this particular event. Compared to the 32,000,000 households in the WECC, 5,000,000 is a high penetration level but not unattainable. At participation levels above the saturation point, the operational burden on any individual device could be lessened by regularly rotating which devices are set to respond first.

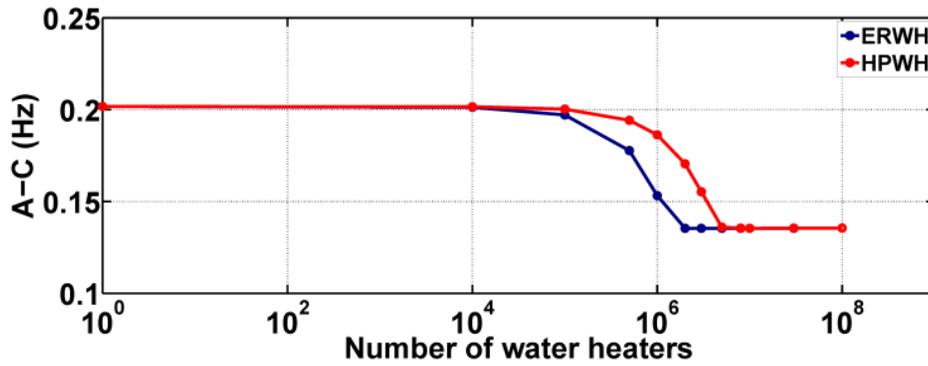


Figure 25. Primary frequency control performance for ERWH and HPWH as a function of participating device population size.

5.0 Conclusions

The system-wide results for a demand-side frequency control strategy were presented that can complement traditional generator controls to maintain the stability of the power system. The gains of the hierarchical load control strategy are determined to meet the NERC BAL-003-1 frequency response standard at both the area level and at the interconnection level. Hardware tests were carried out to test the control logic on physical devices in a laboratory setting and to validate the aggregate load models. The results of detailed simulations of the WECC show that the proposed control strategy can effectively contribute meeting both the system-wide IFRO and area-level FRO, and that the service could be provided by either ERWH or HPWH. Specifically, it was shown that when looking at a particular area in the south, appropriately setting the control gains can lead to more response from controllable loads, thereby contributing towards the FRO area requirement, while adequately meeting the IFRO. A sizeable contribution to the total frequency response required by the WECC could be supplied by electric water heaters alone at penetration levels of less than 15% of the households in the WECC.

6.0 Future Work

Based on the results obtained in FY15 as detailed in this report and feedback received at the Loads as a Resource review meeting in June 2015, the following tasks are proposed for future work.

Investigation of distribution level impacts of using end-use loads for providing primary frequency response: This task would focus on investigating potential issues occurring on distribution feeders, such as voltage fluctuations, conductor/transformer overloads, voltage flicker, etc. with incorporating high penetration of GFAs providing primary frequency response. GridLAB-D, which is an open source distribution software platform, will be used to perform the impact studies. Specifically, the GridLAB-D model will consider an equivalent of the transmission system at the substation to approximate voltage and frequency behavior of the transmission system, using established equivalencing techniques such as Thevenin equivalent. The results of these assumptions will be compared to results of an integrated transmission and distribution simulation environment (GridLAB-D+PowerWorld). The integrated T&D environment can compare results from assumptions of considering an aggregated equivalent of transmission systems in a distribution simulation, and evaluate consideration of aggregated distribution system in a transmission simulation. The effects of these assumptions in both transmission and distribution sides will be investigated.

Measurement & Verification of frequency responsive load control strategies: Measurement and Verification (M&V) is the process of quantifying the quality of a service, typically used to verify that a resource is providing its agreed upon service in a quality manner. To properly incentivize performance, it is necessary to verify that the devices provided the services that they contracted. From the ISO perspective, it is necessary to verify the performance of the aggregator. From the aggregator's perspective, it is necessary to verify the performance of the devices. Currently, primary frequency control is provided by the generators. Systems and rules to verify performance are geared toward generation assets consisting of a small number of large generators with sophisticated communication, control, and monitoring systems with dedicated staff. The M&V methods used for these resources are not scalable when a very large number (~millions) of demand resources participate. The focus of this task will be to develop both supervisory level and device level M&V strategies. The M&V strategies should be sufficiently accurate, while keeping the cost of data acquisition and analysis low in comparison to the value of the service.

Outreach activities: This activity area emphasizes engaging various stakeholders (industry, utilities and DOE) to promote the use of load as a resource for frequency response. The specific organizations we will initially target include CAISO, PJM, ERCOT, and NYISO. The final deliverable will be a transition plan to move from theory and simulations to field testing and finally deployment of the proposed technology.

References

- [1] K. Kalsi, W. Zhang, J. Lian, L. Marinovici, C. Moya, and J. Dagle, "Distributed Smart Grid Asset Control Strategies for Providing Ancillary Services," PNNL-22875, 2013, Pacific Northwest National Laboratory, Richland, WA.
- [2] K. Kalsi, J. Lian, M. Elizondo, W. Zhang, L. Marinovici, and C. Moya, "Loads as a Resource: Frequency Responsive Demand". PNNL-23764, 2014, Pacific Northwest National Laboratory, Richland, WA.
- [3] Standard BAL-003-1 — Frequency Response and Frequency Bias Setting, Version 1, February 7, 2013
- [4] U. Helman, "Resource and transmission planning to achieve a 33% RPS in California-ISO modeling tools and planning framework," in FERC Technical Conference on Planning Models and Software, 2010.
- [5] J. Smith, M. Milligan, E. DeMeo, and B. Parsons, "Utility wind integration and operating impact state of the art," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 900–908, 2007.
- [6] Y. Makarov, C. Loutan, J. Ma, and P. de Mello, "Operational impacts of wind generation on California power systems," *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 1039–1050, 2009.
- [7] A. Molina-Garcia, F. Bouffard and D. S. Kirschen, "Decentralized demand-side contribution to primary frequency control," *IEEE Trans. On Power Systems*, vol. 26, no. 1, pp. 411-419, 2001.
- [8] C. Zhao, U. Topcu, and S. Low, "Frequency-based load control in power systems," in *American Control Conference (ACC), 2012*, 2012, pp. 4423–4430.
- [9] C. Zhao, U. Topcu, and S. Low, "Swing dynamics as primal-dual algorithm for optimal load control," in *Smart Grid Communications (SmartGridComm), 2012 IEEE Third International Conference on*, 2012, pp. 570-575.
- [10] C. Zhao, U. Topcu, and S. Low, "Fast load control with stochastic frequency measurement," *IEEE Power and Energy Society General Meeting*, 2012.
- [11] R. Diao, M. Elizondo, E. Mayhorn and Y. Zhang, "Electric water heater modeling and control strategies for demand response", *Power and Energy Society General Meeting, 2012 IEEE*, 2012, pp. 1–8.
- [12] C. Moya, W. Zhang, J. Lian, and K. Kalsi, "A hierarchical framework for demand-side frequency control," *American Control Conference (ACC), 2014* , vol., no., pp.52,57, 4-6 June 2014.
- [13] M. Elizondo, K. Kalsi, C. Moya, and W. Zhang. 2014. "Frequency Responsive Demand in U.S. Western Power System Model." In *IEEE Power and Energy Society General Meeting* 2015.
- [14] P. Etingov, D. Kosterev, T. Dai, "Frequency Response Analysis Tool," PNNL-23954, 2014, Pacific Northwest National Laboratory, Richland, WA.
- [15] PowerWorld Simulator software, <http://www.powerworld.com>
- [16] Siemens PTI PSS/E software, <https://w3.siemens.com/smartgrid/global/en/products-systems-solutions/software-solutions/planning-data-management-software/planning-simulation/Pages/PSS-E.aspx>
- [17] GE PSLF software, <http://www.geenergyconsulting.com/practice-area/software-products/pslf>
- [18] "WECC MVWG Load Model Report ver.1.0," Western Electricity Coordinating Council, June 12, 2012

- [19] “WECC Composite Load Model (CMPLDW) Specifications,” Western Electricity Coordinating Council, December 4, 2014
- [20] “WECC Composite Load Model with DG Specifications,” Western Electricity Coordinating Council, February 27, 2015
- [21] “Western Electricity Coordinating Council Off-Nominal Frequency Load Shedding Plan”, May 2011
- [22] Widder S, J Peterson, G Parker, and M Beachler. 2013. Demand Response Performance of GE Hybrid Heat Pump Water Heater. PNNL-22642. Pacific Northwest National Laboratory, Richland, WA.
- [23] “Energy Conservation Program: Energy Conservation Standards for Residential Water Heaters, Direct Heating Equipment, and Pool Heaters; Final Rule,” 75 Federal Register 73; April 16, 2010, pp. 20112.
- [24] Mayhorn ET, SA Parker, FS Chassin, RM Pratt. 2015. Evaluation of the Demand Response Performance of Large Capacity Electric Water Heaters. PNNL-23527. Pacific Northwest National Laboratory, Richland, WA.
- [25] ASHRAE. 2006. Method of Testing for Rating Residential Water Heaters. Standard 118.2 -2006. Prepared by the American Society of Heating, Refrigerating, and Air-Conditioning Engineers.
- [26] Lutz, J, A Renaldi, A Lekov, Y Qin, and M Melody. 2012. Typical Hot Water Draw Patterns Based on Field Data. LBNL-4830E. Lawrence Berkeley National Laboratory, Berkeley, CA.
- [27] Akyol B.A., J.N. Haack, S. Ciraci, B.J. Carpenter, M. Vlachopoulou, and C.W. Tews. "VOLTTRON™: An Agent Execution Platform for the Electric Power System." Third International Workshop on Agent Technologies for Energy Systems June 5, 2012, Valencia, Spain.
- [28] B. Larson, M. Logsdon, and D. Baylon, “Residential Heat Pump Water Heater Evaluation: Lab Testing & Energy Use Estimates,” 2011.
- [29] Hendron J and C Engebrecht. 2010. Building America House Simulation Protocols. TP-550-49426. National Renewable Energy Laboratory, Golden, CO.
- [30] CAN/CSA C191-13, 5th Ed.

Distribution

1 Phil Overholt
Department of Energy, OE
1000 Independence Ave., SW
Routing OE-10
Washington, DC 20585

7 **Local Distribution**

Pacific Northwest National Laboratory
Karanjit Kalsi (PDF)
Jianming Lian (PDF)
Tess Williams (PDF)
Laurentiu Marinovici (PDF)
Marcelo Elizondo (PDF)
Jeff Dagle (PDF)
Robert Pratt (PDF)
Jason Fuller (PDF)
Jacob Hansen (PDF)



Pacific Northwest
NATIONAL LABORATORY

*Proudly Operated by **Battelle** Since 1965*

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)

U.S. DEPARTMENT OF
ENERGY

www.pnnl.gov