



Pacific Northwest
NATIONAL LABORATORY

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

Loads as a Resource

Frequency Responsive Demand

December 2015

K Kalsi
J Hansen
J Fuller
LD Marinovici

M Elizondo
T Williams
J Lian
Y Sun

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	M Elizondo
J Hansen	T Williams
J Fuller	J Lian
LD Marinovici	Y Sun

December 2015

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

Pacific Northwest National Laboratory
Richland, Washington 99352

Executive Summary

Current power grid operation predominantly relies on scheduling and regulating generation resources to supply loads and balance load changes. Due to the inherent intermittency of renewable energy, more flexible and fast ramping capacity is required to compensate for the uncertainty and variability introduced by renewable energy resources. With the advancement of information technologies, power system end-use loads are becoming more agile and can participate in the provision of balancing energy and other grid services. The use of demand response can greatly reduce the required generation reserve in a clean and environmentally friendly way.

In FY13, a hierarchical, decentralized control strategy was proposed for thermostatically controlled loads (TCLs) to provide primary frequency response, as discussed in [1]. The proposed control strategy consisted of both supervisory and device layers. In the supervisory layer, a supervisory controller is responsible for gathering system-level information and determining the optimal gains for the TCLs at each bus. In the device layer, individual TCLs switch ON/OFF probabilistically in real time, based on local angle and frequency measurements, so that the aggregated load response under each bus can match the desired power determined by the optimal gains. The optimal controller gains were determined based on decentralized robust control theory, while the switching probabilities of individual TCLs were designed using Markov chain models. In FY14, this control strategy was further extended in [2] so that only local frequency measurements are required and the topology of the power system network can be fully respected. The effectiveness of this approach was demonstrated by large-scale simulation studies on the WECC system using PowerWorld Simulator.

Although the effectiveness of this approach was demonstrated by large-scale simulation studies on the WECC system, it was only applicable to TCLs. Therefore, it was necessary to develop a more generic approach that can control all possible end-use loads, including TCLs and deferrable loads, to provide primary frequency response. The Grid Friendly™ Appliance (GFA) controller, developed at Pacific Northwest National Laboratory, can autonomously switch off various end-use loads by detecting the under-frequency events. In FY13 [1], the impacts of GFAs on the bulk power system frequency stability were investigated, where the GFAs were designed as demonstration units and modeled individually as connected to water heaters. Several important factors regarding the design of the GFA controller and the deployment of GFAs were carefully examined therein. In particular, the performance of the GFAs was evaluated in terms of the response time, geographical location, and penetration level. In FY14, the feasibility of using the GFA controller to provide primary frequency response was investigated in [2]. In particular, the impacts of GFAs on the system frequency response were analyzed by examining the curtailing frequency threshold, which determines the capability of GFAs in providing frequency response. The existing method of selecting the curtailing frequency threshold was insufficient to guarantee a droop-like response from the aggregation of GFAs, due to the inherent frequency deadband. Thus, an improved method of determining curtailing frequency thresholds was then proposed to deal with the existence of the frequency deadband.

In this report, a new frequency responsive load (FRL) controller was proposed based on the GFA controller, which can respond to both over and under-frequency events. A supervisory control was

introduced to coordinate the autonomous response from FRLs in order to overcome the issues of excessive system response due to high penetration of FRLs. The effectiveness of the proposed FRL controller was demonstrated by large-scale simulation studies on the WECC system. Specifically, the FRLs were deployed in the WECC system at different penetration levels to analyze the performance of the proposed strategy, both with and without supervisory level control. While both methods have their own advantages, the case without supervisory control could lead to system-wide instability, depending on the size of the contingency and the number of FRLs deployed in the system. In addition, the voltage impacts of this controller on distribution system were also carefully investigated. Finally, a preliminary measurement and verification approach was also developed.

Acknowledgments

The authors are grateful for the comments and feedback provided by Frank Tuffner at the Pacific Northwest National Laboratory.

Acronyms and Abbreviations

NERC	North American Electric Reliability Corporation
TCL	Thermostatically Controlled Load
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
FRL	Frequency Responsive Load

Table of Contents

Executive Summary.....	iii
Acknowledgments.....	v
Acronyms and Abbreviations.....	vii
Figures.....	x
Tables.....	xii
1 Introduction.....	1
2 Hierarchical demand-side frequency control strategy.....	2
2.1 Frequency Responsive Load Controller.....	2
2.2 Decentralized Threshold Determination.....	4
2.3 Supervised Threshold Determination.....	6
2.4 Measurement and Verification.....	7
2.4.1 Device-level Measurement and Verification.....	7
2.4.2 Supervisor-level Measurement and Verification.....	8
3 System-wide impacts of deployment of improved FRL controllers on the WECC system.....	9
3.1 Modeling Frequency Responsive Loads in PowerWorld.....	9
3.1.1 PowerWorld Integration of Frequency Responsive Loads.....	10
3.1.2 UDM Components for Frequency Responsive Loads.....	11
3.1.3 FRL controllers threshold selection in WECC system model.....	12
3.2 WECC test scenarios.....	12
3.3 Simulation Results from PowerWorld.....	14
4 Investigation of distribution level impacts of using FRLs for providing primary frequency response.....	21
4.1 IEEE 8500-Node Test System.....	21
4.2 Case 0 - Base Scenario.....	23
4.2.1 Distribution System Metrics.....	24
4.3 Case 1 – Supervisory Control of Distributed FRLs.....	26
4.4 Case 2 - Supervisor Control of Distributed FRLs Considering Voltage Impacts of the distribution system.....	31
4.4.1 Case 2-1 Supervisor Control of Distributed FRLs with Voltage Lockout.....	32
4.4.2 Case 2-2 Supervisor Control of Distributed FRLs with Voltage Sorting Approach A.....	36
4.4.3 Case 2-3 Supervisor Control of Distributed FRLs with Voltage Sorting Approach B.....	39
4.4.4 Case 2-4 Supervisor Control of Distributed FRLs with Voltage Lockout and Voltage Sorting Approach A.....	41
5 Conclusions and Future Work.....	45

6	References	47
---	------------------	----

Figures

Figure 1. Flowchart depicting the FRL control logic for frequency events, where the operator ++ indicates that the variable is being counted up by one time step.3

Figure 2. Illustrative examples of the FRL response to different frequency events, where green shaded background illustrates the state of *free*, orange the state of *triggered off*, red the state of *forced off*, and blue the state of *released off*.4

Figure 3 Illustrative example of improved threshold determination5

Figure 4 Diagram of hierarchical decentralized control strategy6

Figure 5 Illustration of supervised frequency threshold determination, where $f_2 < 60$ for the ON group, and $f_2 > 60$ for the OFF group7

Figure 6. Diagram of user-defined model (UDM) for aggregate controllable load in PowerWorld and its interface with the WECC composite load model [14].....10

Figure 7. WECC balancing authorities (right) and sub regions for reserve sharing groups (left) [16]13

Figure 8. Low availability scenario. (a) Frequency at one representative bus with controllable loads. (b) Controllable load value at the bus.15

Figure 9. Low availability scenario - System frequency response in MW/Hz16

Figure 10. High availability scenario. (a) Frequency at one representative bus with controllable loads. (b) Controllable load value at the bus.17

Figure 11. High availability scenario - System frequency response in MW/Hz17

Figure 12. Extreme scenario. (a) Frequency at one representative bus with controllable loads. (b) Controllable load value at the bus.18

Figure 13. Extreme availability scenario - System frequency response in MW/Hz18

Figure 14. Frequency responses. (a) Decentralized frequency threshold determination. (b) Supervised frequency threshold determination.....19

Figure 15. Decentralized vs. supervised threshold determination. (a) Low availability. (b) High availability. (c) Extreme availability.20

Figure 16 Load profile for the IEEE 8500 node system.23

Figure 17 Number of water heaters that are on.....24

Figure 18 Voltage profile for the base case taken at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks.25

Figure 19 Under-frequency event from tripping two generator in the south of WECC (~2.7 GW)..27

Figure 20 Response in MW for Case 1 with under-frequency event.27

Figure 21 Voltage profile for Case 1 with under frequency event taken at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.28

Figure 22 Over-frequency event from tripping large load in the south of WECC (~0.9 GW)29

Figure 23 Response in MW for Case 1 with over frequency event.....29

Figure 24 Voltage profile for Case 1 with over-frequency event taken at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.30

Figure 25 Response in MW for Case 2-1 with the under-frequency event.32

Figure 26 Voltage profile for Case 2-1 with the under-frequency event taken at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.33

Figure 27 Response in MW for Case 2-1 with the over-frequency event.....34

Figure 28 Voltage profile for Case 2-1 with the over-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.34

Figure 29 Tap changes for VREG3 (phase A, B, and C) along with switching operations for CAPBANK1 (phase B) during the over-frequency event with voltage lockout enabled.36

Figure 30 Voltage profile for Case 2-2 with the under-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.37

Figure 31 Voltage profile for Case 2-2 with the over-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event. The above conclusion is further supported by the violation summary in Table 12.38

Figure 32 Voltage profile for Case 2-3 with the under-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.39

Figure 33 Voltage profile for Case 2-3 with the over-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.41

Figure 34 Response in MW for Case 2-4 with the under-frequency event.42

Figure 35 voltage profile for Case 2-4 with under frequency event taken at 14:31. Red dots corresponds to voltage regulators and blue dots are capacitor banks. '+' symbols corresponds to water heater changing state due to the emergency event.....42

Figure 36 Response in MW for Case 2-4 with over-frequency event.43

Figure 37 Voltage profile for Case 2-4 with over-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.....44

Tables

Table 1 Regulator and capacitor settings (modified from original system).....	22
Table 2 Voltage violations for base case.....	25
Table 3 Transformer and line overloads	26
Table 4 Number of voltage regulator tap changes and capacitor bank switching	26
Table 5 Violation data for Case 1 during under-frequency event.	28
Table 6 Violation data for Case 1 during over-frequency event.....	30
Table 7 Capacitor switches and regulator tap changes for Case 1 during the over-frequency event.	31
Table 8 Violation data for Case 2-1 during the under-frequency event.....	33
Table 9 Violation data for Case 2-1 during the over-frequency event.	35
Table 10 Regulator tap changes and capacitor switches for Case 2-1 during the over-frequency event.	35
Table 11 Violation data for Case 2-2 during the under-frequency event.....	37
Table 12 Violation data for Case 2-2 during the over-frequency event.	38
Table 13 Violation data for Case 2-3 during the over-frequency event.	40
Table 14 Violation data for Case 2-3 during the over-frequency event.	41
Table 15 Violation data for Case 2-4 during under frequency event.	43
Table 16 Violation data for Case 2-4 during the over-frequency event.	44

1 Introduction

The vast integration of renewable energy into the electric power grid imposes daunting challenges to the traditional centralized management system. On the one hand, the variability and uncertainty of renewable generation will substantially increase the need for operational reserves to balance supply and demand instantaneously and continuously, as pointed out in [4] and [5]. On the other hand, the total system inertia, as well as contingency reserve, decreases as non-dispatchable renewable generation displaces conventional generation. Therefore, it becomes extremely difficult for the system operator to maintain the stability and reliability of the power grid. Hence, today's renewable penetration is still limited due to the lack of the technologies that are able to reliably and affordably manage the dynamic variability introduced

In order to maximize renewable penetration for energy efficiency in the future, more flexible and fast ramping capacities are needed to handle the variability of such uncontrollable resources. End-use loads such as air conditioners, water heaters, washers, and dryers can increasingly be actively controlled to provide grid services. An aggregation of such demand-side assets can be properly coordinated with supply-side generation to provide operating reserve for system needs. From a reliability perspective, reserves can be provided by either generation or demand side resources, as long as they respond with the speed, accuracy and magnitude that is required. For example, a large population of end-use loads have a fast aggregate ramping rate, and thus present an enormous potential to mitigate the uncertainty from renewable generation. The coordination between generation and flexible loads could reduce the need to build new transmission facilities to accommodate large amounts of renewable generation.

In FY15, the original Grid Friendly™ Appliance (GFA) control logic was first extended to develop a new Frequency Responsive Load (FRL) controller that can autonomously respond to both under- and over-frequency events. Then, a simple decentralized control strategy was derived based on the improved design proposed in [2]. Although it enabled autonomous response, which is critical for frequency response, the effectiveness of the decentralized control strategy is limited due to several practical factors. One of the significant factors is the penetration level of FRLs in the system. When there is high penetration, it is inevitable to have excessive response from FRLs, which could potentially negatively impact system stability. Hence a supervisory control was introduced to appropriately coordinate the autonomous response from FRLs. The developed hierarchical decentralized control strategy consists of two decision making layers, including supervisory and device layers. The coordinator at the supervisory layer coordinates autonomous response to ensure the aggregated response to be droop-like without the issue of becoming excessive due to high penetration. The proposed control strategy also preserves the autonomous operation of the end-use loads.

In this report, the effects of the proposed control strategy on the distribution system is investigated. Distribution utilities have the responsibility for maintaining service and reliability at this level. When resources are being controlled, and particularly synchronized, outside of the control of the local reliability coordinator, there is the potential for causing unforeseen issues on the distribution system. Initial results for both under and over frequency events are reported. Furthermore, voltage protection algorithms are developed and their effectiveness evaluated.

2 Hierarchical demand-side frequency control strategy

The GFA controller monitors system frequency locally and turns off the appliances under control when under-frequency events are detected [3]. The feasibility of using the GFA controller to provide primary frequency response has been demonstrated in [6]. However, the GFA only responds to under-frequency events. In order to ensure the applicability of load control strategy over a wide range of operating conditions, it is desirable that end-use loads respond to both under- and over-frequency events. Therefore, a new controller for FRLs is proposed that extends the functionality of the GFA controller.

2.1 Frequency Responsive Load Controller

A detailed flowchart showing the control logic of the proposed FRL controller is depicted in Figure 1. Functionally, each individual FRL has two different operating modes: *under-frequency* ($f \leq 60$ Hz) and *over-frequency* ($f > 60$ Hz). At any given time, the FRL can only be operated in one mode, which is determined according to the local frequency measurement. The two operating modes can be further divided into seven different states: *free*, *triggered off*, *triggered on*, *forced off*, *forced on*, *released off*, and *released on*. In the *free* state, the FRL evolves based on its internal dynamics, turning ON or OFF according to its internal control. Once the FRL controller detects that the grid frequency falls below a predetermined curtailing frequency threshold f_t^u (or above a predetermined rising frequency threshold f_t^o), the FRL changes its operating state from *free* to *triggered off* (or *triggered on*). It remains in this state as long as the grid frequency does not return above f_r^u (or go below f_r^o). If the frequency event persists longer than the response time $T_{b_t}^u$ (or $T_{b_t}^o$), the FRL shuts down the device and switches it from *triggered off* (or *triggered on*) to *forced off* (or *forced on*). The time periods $T_{b_t}^u$ and $T_{b_t}^o$ are defined by the low-pass digital filter that smooths the frequency measurements to avoid reactions to unreliable data or noise. Once the grid frequency rises above a predetermined restoring frequency threshold f_r^u (or falls below a predetermined restoring frequency threshold f_r^o), where $f_r^u > f_t^u$ (or $f_r^o < f_t^o$), the FRL switches from *forced off* (or *forced on*) to *released off* (or *released on*). The FRL remains in its state if the frequency stays above f_r^u (or below f_r^o). If the FRL has been in the state of *released off* (or *released on*) for longer than the release time delay $T_{b_r}^u$ (or $T_{b_r}^o$), the FRL switches its state back to *free*, and follows its internal dynamics. The release time delays $T_{b_r}^u$ and $T_{b_r}^o$ are designed for the purpose of preventing the rebound effect that occurs when all the FRLs try to return to normal operations at the same time.

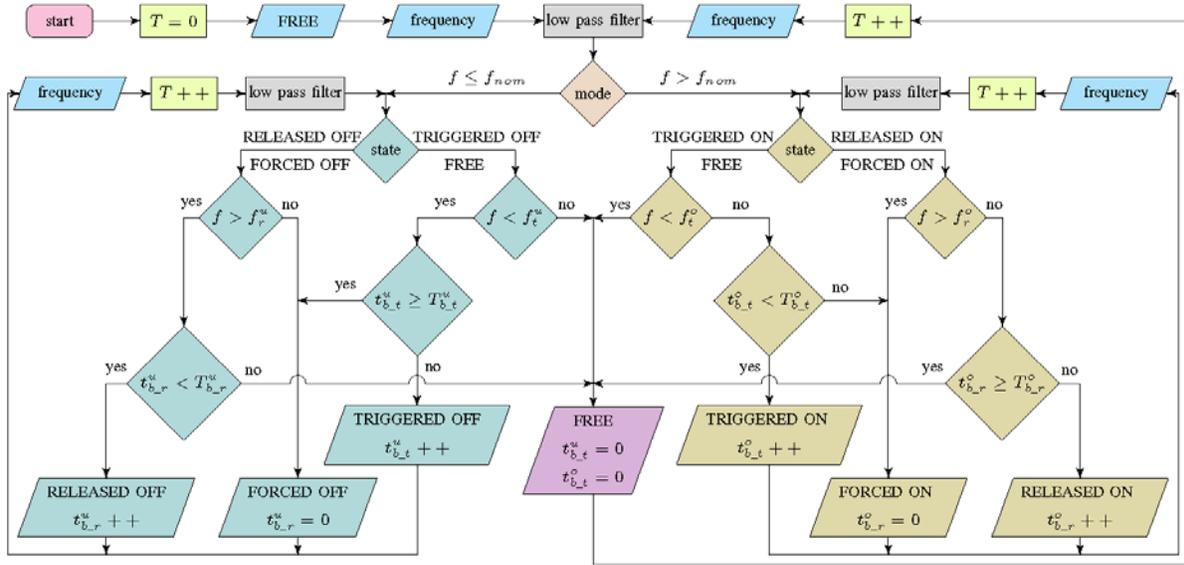


Figure 1. Flowchart depicting the FRL control logic for frequency events, where the operator ++ indicates that the variable is being counted up by one time step.

The response of the FRL controller to various frequency events is further illustrated by two examples shown in Figure 2. In these two examples, we only consider the case of under-frequency events. Similar examples can be constructed for the case of over-frequency events. In the top plot, the FRL starts out in the state of *free* (green) when the frequency starts to dip. When the frequency drops below the curtailing frequency threshold f_t^u , the FRL changes its state to *triggered off* (orange). Then, the frequency is restored above the restoring frequency threshold f_r^u within the response time $T_{b,t}^u$, so the FRL changes its state back to *free* (green), resuming normal operation.

In the bottom plot of Figure 2, the FRL starts out in the state of *free* (green) as well. When the frequency drops below the frequency threshold f_t^u , the FRL changes its state to *triggered off* (orange). In this case, the frequency is not restored above the frequency threshold f_t^u within the response time $T_{b,t}^u$, so the FRL changes its state to *forced off* (red). The FRL stays in the state of *forced off* until the frequency is restored above the frequency threshold f_r^u , and then changes its state to *released off* (blue). However, the frequency does not stay above f_r^u for enough time, so the FRL changes its state back to *forced off* (red). After some time, the frequency returns above f_r^u again and the FRL changes its state to *released off* (blue) once time. Finally, the frequency stays above the f_r^u for a longer time than the release time $T_{b,t}^u$, so the FRL changes its state to *free* (green) resuming the normal operation.

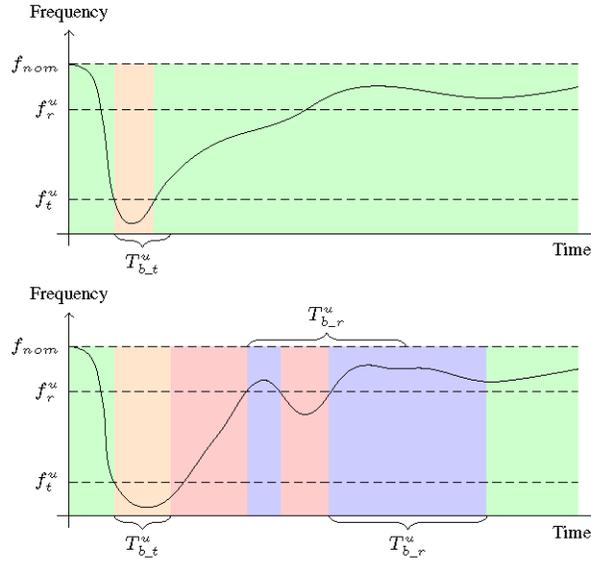


Figure 2. Illustrative examples of the FRL response to different frequency events, where green shaded background illustrates the state of *free*, orange the state of *triggered off*, red the state of *forced off*, and blue the state of *released off*.

2.2 Decentralized Threshold Determination

It is highly desirable to have the primary frequency response provided from the demand side to be compatible with that from the supply side, where generator power changes proportionally to frequency deviation according to a droop curve. Hence, in the proposed control, the curtailing and rising frequency thresholds are determined such that the aggregated response of engaged FRLs during frequency events is droop-like.

In the original design of the GFA controller, the curtailing frequency thresholds were randomly selected from a prescribed range based on a uniform distribution by individual GFAs in a decentralized way. For a field demonstration, this frequency range was chosen to be between 59.95 Hz and 59.985 Hz in [7], based on analysis of historical frequency data. As pointed out in [6], this original threshold determination is insufficient to guarantee the desired performance. This is because the existence of frequency deadband prevents the aggregated response of online GFAs from being droop-like, especially when the frequency deviation is shallow. In order to deal with the restriction imposed by the necessity of the frequency deadband, it is proposed in [6] to select the curtailing frequency thresholds in a different way. This ensures the droop-like response, even in the presence of the frequency deadband. According to the improved method, individual GFAs should first randomly pick their curtailing frequency thresholds from the prescribed range between 59.95 Hz and 60 Hz based on a uniform distribution. If the selected threshold is between 59.985 Hz and 60 Hz, it will be reset to 59.985 Hz. The resulting distribution of curtailing frequency thresholds is illustrated in Figure 3(a), which shows 1000 online GFAs with power uniformly distributed between 4 kW and 6 kW. The corresponding relationship between the power reduction and the frequency deviation at any time instant is described by a Power versus Frequency (PF) curve in Figure 3(b). With the improved threshold determination, the aggregated

response of online GFAs can well approximate the desired droop-like response for small frequency deviations in spite of the existence of frequency deadband.

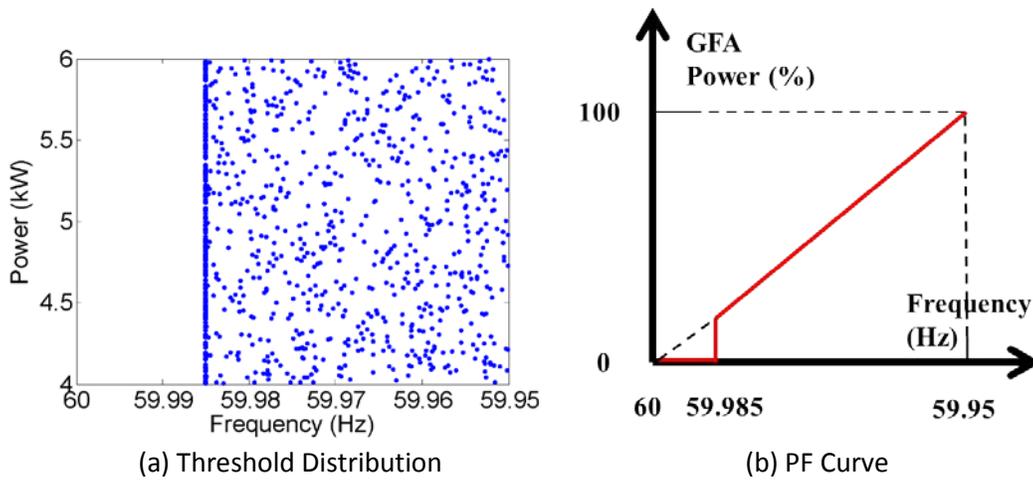


Figure 3 Illustrative example of improved threshold determination

By applying the improved method of threshold determination proposed in [6] to determine both curtailing and rising frequency thresholds, we can easily derive a decentralized control strategy for engaging end-use loads to provide primary frequency response using the proposed FRL controller. The decentralized control strategy enables autonomous response from end-use loads during frequency events, which is critical for primary frequency response. However, there are several practical factors to prevent the decentralized control strategy from providing the desired response. First, power ratings of engaged FRLs may be non-uniformly distributed. In order to achieve an approximate droop-like PF curve as shown in Figure 3(b), it is important for the power ratings of engaged FRLs to be uniformly distributed. In the worst case, the resulting PF curve may be so biased that a large amount of FRLs will respond to small frequency deviations. In this case, the excessive response from FRLs will negatively impact the overall system response, which could potentially render the system unstable. Second, even when droop-like response from FRLs is achieved, it will be inevitably excessive when the penetration level of FRLs is high. This is because the range of curtailing and rising frequency thresholds is predetermined and independent of the amount of engaged FRLs at different points in time.

2.3 Supervised Threshold Determination

In order to overcome the identified drawbacks, but maintain the advantages of the decentralized control strategy, a hierarchical decentralized control strategy as shown in Figure 4 is proposed. The control strategy consists of two decision making layers: supervisory and device. In the supervisory layer, the coordinator is responsible for ensuring the aggregated response from engaged FRLs to be droop-like during frequency events and preventing it from being excessive under high penetration of FRLs. In the device layer, the autonomous and rapid response to frequency contingency events is implemented as described in Section 2.1.

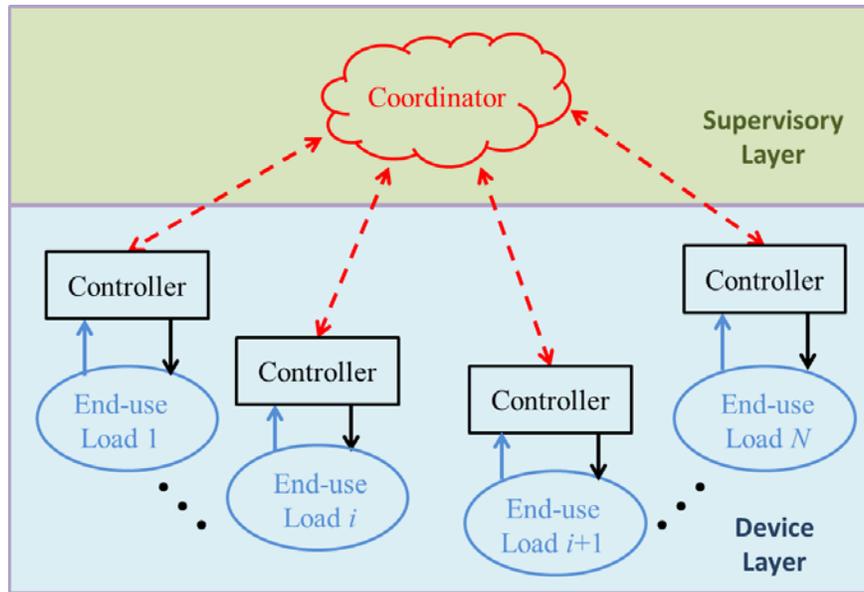


Figure 4 Diagram of hierarchical decentralized control strategy

Under the proposed strategy, the FRLs communicate with the coordinator once every control period. The length of the control period is a design parameter dependent on the characteristics of supervised FRLs. Further research is required to determine the appropriate length of the control period. At the beginning of each control period, the FRLs in the free mode submit information including power rating (kW) and power mode (ON or OFF) to the coordinator. After collecting all the information, the coordinator first divides the engaged FRLs into two groups. The ON group consists of those FRLs that are currently ON, and will provide under-frequency response. The OFF group consists of those that are OFF, and will provide over-frequency response. Then the coordinator calculates the total aggregated power of each group, P_{\max} and selects the desired droop value R for each group, based on the corresponding magnitude of P_{\max} . Once P_{\max} and R are determined, the boundary frequency f_2 for the range of frequency thresholds is automatically fixed for each group, as shown in Figure 5, where the other boundary frequency f_1 is selected for the frequency deadband of each group. Finally, the coordinator assigns and broadcasts the frequency thresholds to each group of FRLs so that the droop-like response during frequency events can be ensured. After receiving the frequency thresholds from the coordinator, individual FRLs update their own controller settings. In real time, they will monitor the grid frequency locally, and respond to the detected frequency events independently. By determining f_2 indirectly through the selection of R , the maximum frequency deviation to be responded to becomes dependent

on the penetration level of FRLs, which effectively overcomes the issue of excessive response under high penetration of FRLs.

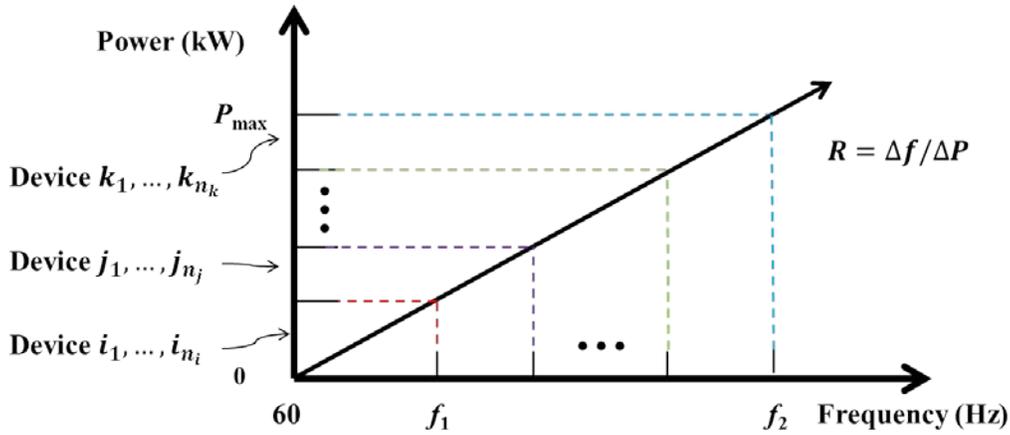


Figure 5 Illustration of supervised frequency threshold determination, where $f_2 < 60$ for the ON group, and $f_2 > 60$ for the OFF group

2.4 Measurement and Verification

Measurement and Verification (M&V) is the process to check and verify that a product or service is meeting the expected quality and the needs of users. The M&V procedure is independent of any model assumptions and should be only depending on measurements. M&V procedures are commonly used for physical devices, such as generators. It has also been used for evaluating the aggregate performance of a population of small devices, such as residential air conditioners and water heaters [8]. The M&V task for the proposed frequency response load control strategies considers two levels: supervisor level M&V and device level. For the supervisor-level M&V, it needs to verify that the population is providing a droop-like response as described in Section 2. For the device-level M&V, it needs to verify that each individual device is providing the promised response. In this section, we will describe the general procedure to implement M&V for both levels.

2.4.1 Device-level Measurement and Verification

At the device level, the purpose of M&V is to check whether an engaged device is providing the promised response during contingencies. That is, when the measured frequency is consistently below the local curtailing frequency threshold for sufficient amount of time, the device that is currently ON will definitely turn OFF, as specified by the local controller. Hence, there are two things to verify at the device level. First, it needs to verify that the device can detect the frequency events locally. Second, it needs to verify that the device can change its operating mode whenever the measured frequency drops below the curtailing frequency threshold and remains there.

In order to complete the device-level M&V, the device needs to maintain two logs at the time resolution of second. One log is to record the measured frequency from power outlet. This log will be used to verify whether the device can detect the frequency event correctly. The other log is to record the operating mode of the device. This log will be used to verify whether the local controller can turn off or turn on the device at the right time. These two logs will be submitted to the supervisor within a given amount of time (for example, 7 days) after each identified frequency event.

2.4.2 Supervisor-level Measurement and Verification

At the supervisor level, the purpose of M&V is to verify whether the actual aggregated response of engaged FRLs follows the expected droop-like response, as specified by the desired droop value of R . Such verification will be performed for each control period whenever there are frequency events. The supervisor collects logs of both local frequency measurements and operating modes from individual devices and then calculates the ratio of the actual aggregated response of engaged FRLs to the expected aggregated response. Here, we assume that the supervisor will keep track of the engagement of individual devices over each control period, and also the corresponding power consumption of individual devices.

Although it is technically feasible to calculate the actual aggregated response, it is cost-prohibitive if we request all the devices in the device layer to submit their logs. This is because the number of end-use loads is usually so large that the resulting burden of communication and computation will be very high. In order to make the cost of M&V smaller in comparison to the value of the provided service, we can only consider performing the verification with a subset of the population instead of the whole population. Therefore, the selected subset of the population has to be carefully planned so that the survey of the small set can represent the whole population as accurately as possible. The detailed sampling plan will be described in the rest of this section.

We assume that the supervisor has the information about the power capacity of individual devices. This information can be collected at the time when customers sign up for the demand response programs. Based on the collected information, the distribution of the power capacity for the whole population can be obtained. The selected set of participating devices has to be randomly chosen from the whole population in such a way that the distribution of the power capacity for the selected subset should be similar to the overall distribution. This can be done by categorizing the whole population into different bins based on individual power capacities. The number of bins needs to be sufficient to reflect the shape of the overall distribution. For example, if the distribution is close to a uniform distribution, the number of bins can be small and with the same bin length; if the distribution is multimodal, different bins should be assigned for different parts of the distribution such as peak and off-peak areas with various bin lengths. Once the bins for the power capacities are determined, a portion of the population within each bin is randomly selected. The percentage of the selected customers in each bin needs to be consistent for all bins.

After the subset of the population is determined, the devices within the subset will be used for both device-level and supervisor-level M&V. They will maintain the logs of frequency measurement and operating modes, and then submit these logs to the supervisor with a given amount of time after each

frequency event. Then the supervisor will utilize the collected logs to calculate the ratio of the actual aggregated response to the expected aggregated response during frequency events. Each supervisor has to maintain this ratio to be within a predefined threshold. The subset of the population needs to be resampled after a period of time (e.g., after several contingencies have occurred) with the same sampling criteria. Eventually, all participating customers in the program will be involved in one or more M&V activities.

3 System-wide impacts of deployment of improved FRL controllers on the WECC system

In this section, the improved FRL described in Section 2 is implemented in a detailed model of the Western Electricity Coordinating Council (WECC). The system-wide impacts of the FRLs at varying levels of deployment on the frequency response to a contingency are simulated using the PowerWorld Simulator. The performance of the FRLs coordinated by a supervisor is contrasted with autonomous FRLs.

3.1 Modeling Frequency Responsive Loads in PowerWorld

Utility and research engineers typically perform large power system interconnection studies with commercial software, like Siemens PTI PSS/E [11], GE PSLF [12], and more recently, the PowerWorld Simulator [10]. 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 UDMs until the model becomes more widely used, after which the models can be incorporated into standard libraries or as a common feature. This section describes the implementation of the model used for the hierarchical demand-side frequency control on the WECC interconnection model.

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. Such models can either be static, such as the ZIP model (load modeled as a combination of constant impedance, constant current, and constant power components) or dynamic, such as motor load models. Currently, load models have been implemented in the WECC system model as a composite load model, as described in [13]-[15]. 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 [13]-[15].

3.1.1 PowerWorld Integration of Frequency Responsive Loads

One of the components of the composite load model in PowerWorld is the static load, which encompasses electric resistance water heaters. A UDM was created to model a population of water heaters and interfaced with the WECC composite load model through the static load component. Figure

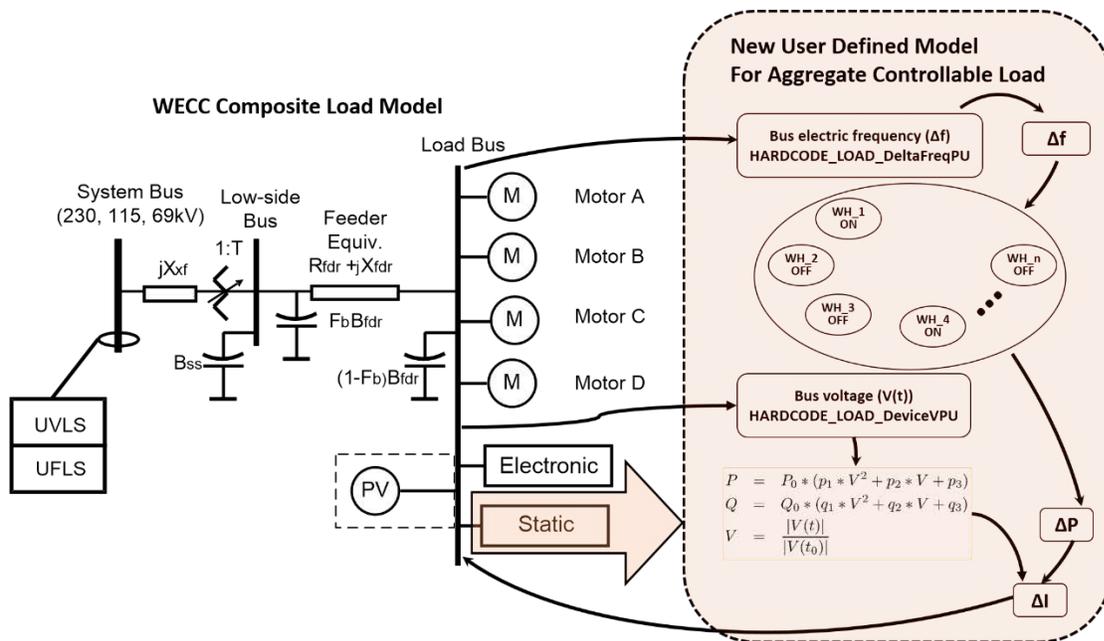


Figure 6. Diagram of user-defined model (UDM) for aggregate controllable load in PowerWorld and its interface with the WECC composite load model [14]

6 shows a diagram of the UDM in the WECC composite load model.

The UDM interacts with the network simulator through the values of bus frequency for control and voltage for load modeling at each simulation time. The UDM uses the updated bus frequency and voltage to model load dynamics representing the natural state transition of water heaters, the hierarchical demand-side frequency control, if activated, and the natural load dependence on voltage. The updated power consumption of the population of water heaters is communicated to the rest of the power system model at each time step of simulation in the form of current injection.

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. The UDM consists of two levels of abstractness, as depicted in Figure 1. The first is the individual water heater behavior according to the Equivalent Thermal Parameter (ETP) model and, if purposely set, controlled by the deviation of frequency Δf from the nominal value. The second is the aggregate load behavior of the currently on devices as simple ZIP load model that captures the voltage dependence of the resistive loads.

To ensure correct initialization of the load UDM, the bus steady state solution returned by the simulator power flow solver is used to initialize parameters into the load model. Part of the initial aggregate load is given by a steady state population of water heaters in the ON state. The transitions of the water heaters in time are calculated by:

$$P_{UDM}(t) = P_{UDM}^{SS} + \Delta P_{WH}(t)$$

$$\Delta P_{WH}(t) = P_{WH}^{ON}(t) - P_{initWH}^{ON}$$

with

- P_{UDM}^{SS} – aggregated power at the bus where water heaters are connected at $t = 0$,
- $P_{UDM}(t)$ – aggregated power at the bus where water heaters are connected as a function of time,
- P_{initWH}^{ON} – power of the water heaters in the ON state at $t = 0$,
- $P_{WH}^{ON}(t)$ – power of the water heaters in the ON state as a function of time.

3.1.2 UDM Components for Frequency Responsive Loads

At the level of each UDM, the load thermal dynamic behavior is captured by a particular population of water heaters, where each individual water heater has a thermal model with distinct parameters, as described by the one-node model below:

$$T(t + \Delta t) = T(t)e^{a\Delta t} + \frac{b}{a}(e^{a\Delta t} - 1)$$

where

$$a = \frac{-\dot{m}C_p - UA}{C_w}$$

$$b = \frac{\dot{m}C_p T_{in} + UA T_{amb} + Q_{elec}}{C_w}$$

with

- $T(t)$ – water temperature in the tank at time t
- Δt – sampling time interval
- \dot{m} – inlet water mass flow
- C_p – water heat capacity

C_w – water thermal capacitance
 T_{in} – inlet water temperature
 T_{amb} – ambient temperature
 UA – conductance to ambient conditions
 Q_{elec} – water heater heat input rate (rated power of water heater)

A diverse population of water heaters is created by randomly assigning heterogeneous parameters in a uniform distribution with support given by commonly used minimum and maximum values. For each ETP-described device, the FRL controller monitors the frequency at the device power source and forces a state transition if the local frequency decreases below a threshold, which affects the natural transitions of the water heaters. For the system used in this study, it is assumed that 17% of water heaters are initially in the ON state and about 10 daily natural state transitions occur per device.

3.1.3 FRL controllers threshold selection in WECC system model

The frequency-based controller at each water heater monitors the frequency at the device power source. Based on a device-independent threshold, the controller will turn off the device in the case of an under-frequency event until the system frequency is restored to a satisfactory value. This section considers the impact of varying penetrations of FRLs in the WECC and contrasts how decentralized versus supervised methods of assigning curtailment thresholds to water heaters could potentially improve or worsen the overall system behavior during and in the seconds following an under-frequency event. As presented in Section 2.2, the decentralized threshold assignment is carried out during manufacturing. Water heater manufacturers allocate thresholds between 59.5 Hz and 59.985 Hz to the FRL controllers in an independent and random manner. In the supervised control strategy, as detailed in Section 2.3, a coordinator assigns thresholds to available loads at fixed time intervals based on the desired frequency response that depends on the amplitude of the contingency. The autonomous frequency response of the water heaters in response to a contingency is implemented in the same manner in the two methods; the only difference is the manner of assigning frequency thresholds.

3.2 WECC test scenarios

The WECC 2015 heavy-load summer case was used to test the hierarchical demand-side frequency control connected to the WECC composite load model. The heavy-load summer case was also used to show how the controller responds to a contingency to maintain system stability. This subsection provides a summary of the WECC system model and the scenarios of controllable loads in terms of size of participating load and frequency assignment strategy.

The WECC system has 20,573 buses, 3,984 generators, 16,426 transmission lines, and 10,922 loads. As shown in the left-hand panel of Figure 2, the WECC system covers the western parts of the United States, Canada, and the northern portion of Baja California, Mexico. The 2015 heavy-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-hand panel of Figure 7. The balancing authorities are 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 authority, substation, and individual device. The desired behavior of FRLs could be defined at each level of coordination.

In this report, the coordination is illustrated, showing how controllable loads can supplement the droop response provided by generators, quantified by the MW/Hz response. The requirement for WECC is a load drop of 8400 MW per 1.0 Hz deviation from nominal at settling frequency [17]. The WECC model utilized in this example exceeds this requirement, as will be shown in Section 3.3. Instead of focusing on the exact 8400 MW/Hz of the requirement, the study discussed in this report shows how a

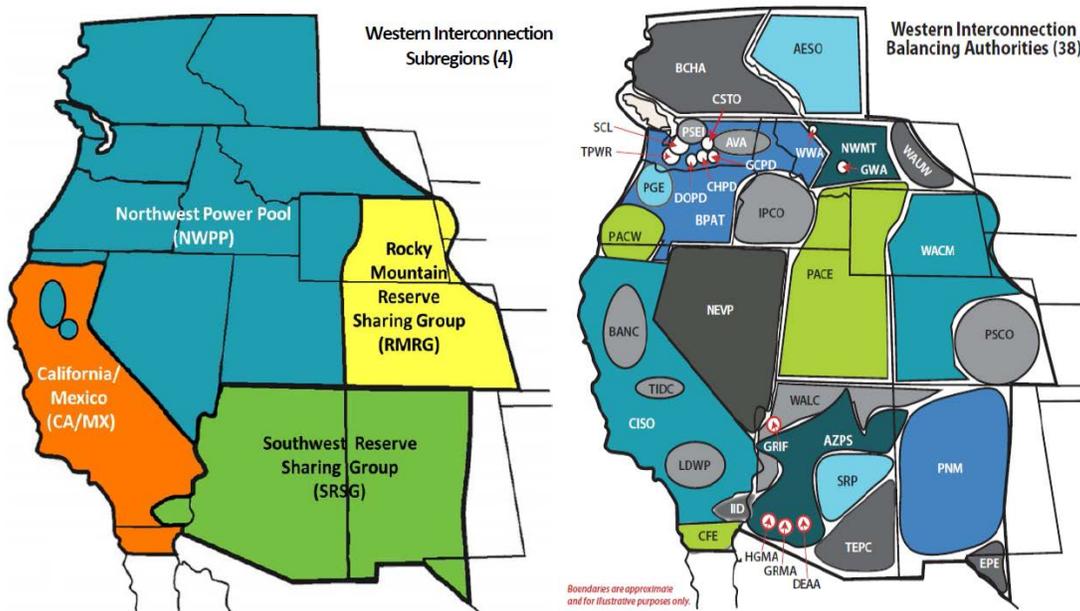


Figure 7. WECC balancing authorities (right) and sub regions for reserve sharing groups (left) [16]

specific additional MW/Hz droop-like response can be provided by controllable loads across the WECC system. Moreover, the study discusses the pros and cons of a decentralized versus a supervised strategy for assigning curtailment frequency thresholds.

Based on the level of controllable load penetration, three main scenarios were simulated:

1. Scenarios with low availability of controllable load, about 900,000 water heaters (about 700MW in ON state)
2. Scenarios with high availability of controllable load, about 6.2M water heaters (about 4.6GW in ON state)
3. Scenarios with extreme availability of controllable load, about 13M water heaters (about 10GW in ON state)

Each of these scenarios included three sub-cases, according to the control strategy:

- A. Water heaters follow their natural transitions due to water draw and tank water temperature dynamics, not monitoring or responding to frequency
- B. Water heaters follow their natural behavior, while monitoring and responding to frequency deviation according to fixed thresholds for frequency response; predefined curtailment thresholds are set based on a uniform distribution between 59.5 Hz and 59.985 Hz
- C. Water heaters behave naturally, while monitoring and responding to frequency deviation according to thresholds coordinated by a supervisor; the supervisor distributes curtailment thresholds based on the desired load contribution to the system frequency response, in order to achieve a certain drop in load for a given change in frequency.

In each of the scenarios, the FRLs have an average rated power of 4.5 kW and are uniformly geographically distributed to 5,723 instances of UDMs, and connected to the WECC model at 5,723 buses according to Figure 6 and the details in Subsection 3.1.1. The FRLs are desired to provide an additional 7,955MW/Hz (5% of 159,104 MW total steady state load of the WECC system) frequency response. This value represents an expected response for the entire system, complementing the response from generators. The specific value is adopted for illustrative purposes, but in practice, a value could be assigned, for example, to complement contribution from generation to the requirements of NERC standards [17].

3.3 Simulation Results from PowerWorld

This subsection discusses the results of the scenarios enumerated in Section 3.2, in order to demonstrate the efficacy of the hierarchical demand-side frequency controller in trying to contribute a certain amount of frequency response to the whole system. In all cases, the system response was observed as it recovers from an under-frequency event caused by losing two major generators, resulting in a total power loss of 2,756 MW. In the first scenario, there is a low number of frequency-controlled water heaters available to participate in the hierarchical demand-side frequency control. Out of the total of 159,104 MW of load in the WECC, only 700 MW could assist in arresting the frequency drop due to an under-frequency event. The 700 MW is provided by 900,000 water heaters distributed throughout the WECC, with a given probability of being found in the ON state at 17% and an average rated power of 4.5 kW.

Analyzing the frequency at a representative bus where controllable loads are located, the results in Figure 8(a) show that the frequency deviation was arrested faster in both S1-B and S1-C, as compared with S1-A where end-use devices do not contribute to frequency response. Moreover, after the transient, the frequency settles to values closer to the nominal. The overall system recovery is

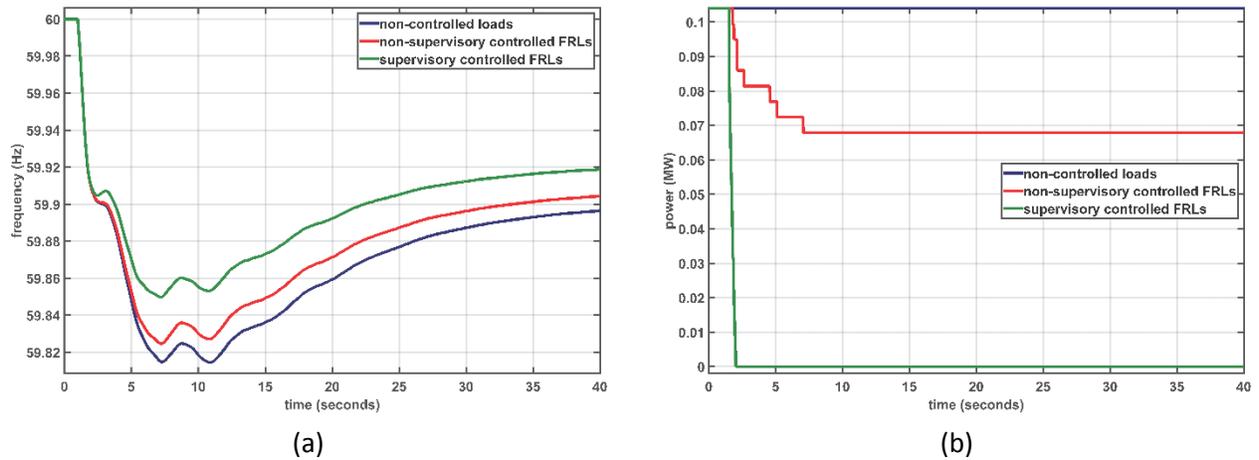


Figure 8. Low availability scenario. (a) Frequency at one representative bus with controllable loads. (b) Controllable load value at the bus.

accomplished by dropping certain amount of load at each bus with frequency responsive loads based on their threshold values, as shown in Figure 8(b).

The difference in the frequency after it reaches steady state in scenarios S1-B and S1-C suggest that the decentralized threshold determination was not able to set the curtailment frequencies such that enough load changes to OFF state when needed to provide the desired frequency response. On the contrary, by taking into account the desired overall load contribution to the system's frequency response, the supervised threshold determination method distributed the curtailment threshold such that enough load would be dropped to reach a value for the frequency response closer to the intended one. According to Figure 9, at 40 seconds, when frequency has stabilized after the contingency, the total system frequency response (the contribution from both generators and loads) is roughly 28,859 MW/Hz for scenario S1-B and 33,940 MW/Hz for scenario S1-C, respectively. Relative to scenario S1-A when loads are not frequency responsive, and the frequency response of about 26,654 MW/Hz is due to generators only, there was an increase of 2,205 MW/Hz for S1-B and 7,286 MW/Hz for S1-C. Scenario S1-B failed to get close to the desired load drop of 7,955 MW/Hz because of the low number of available frequency-controlled devices, and curtailment thresholds being uncorrelated with this value. Scenario

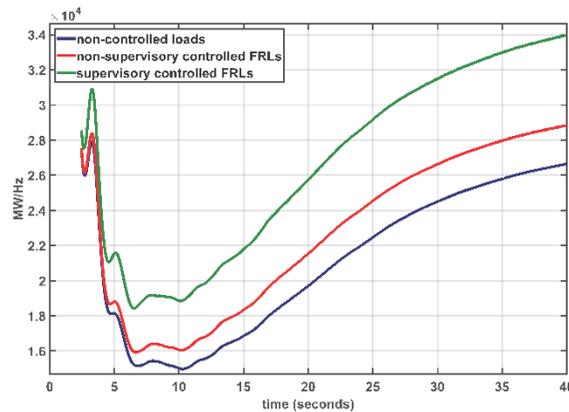


Figure 9. Low availability scenario - System frequency response in MW/Hz

S1-C, however, manages to reach of value very close to the desired one, supporting the conclusion that a supervised threshold determination is of great help when in need to drop enough load to assist with the frequency response, but not much is available.

In the second scenario, WECC has approximately 6.2M FRLs that are uniformly distributed across 5,723 buses. At an average of 4.5 kW per water heater and with a 17% of them being initially ON, that results in about 4.6GW of available controllable loads. Similar to the first scenario, the frequency deviation is arrested faster when FRLs are involved in frequency control. However, as seen in Figure

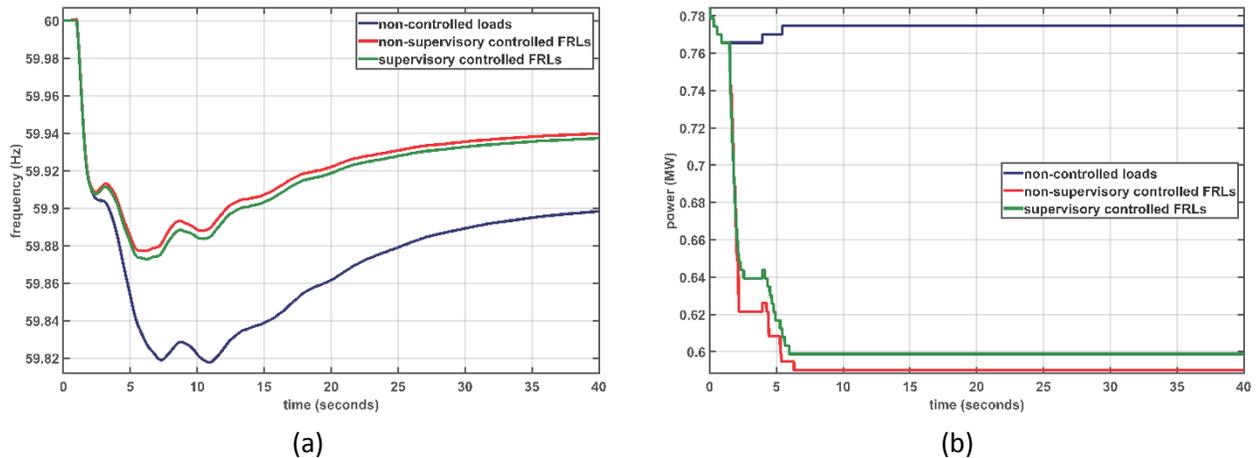


Figure 10. High availability scenario. (a) Frequency at one representative bus with controllable loads. (b) Controllable load value at the bus.

10(b), both the non-supervised and the supervised threshold determination lead to almost the same amount of load to be dropped at each bus equipped with FRLs. Hence, both cases have roughly identical frequency recovery, as seen in Figure 10(a).

This suggests that the frequency responses are also close to each other for S2-B and S2-C. Figure 11

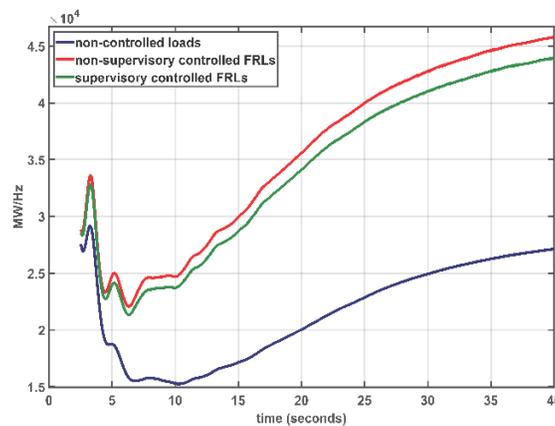


Figure 11. High availability scenario - System frequency response in MW/Hz

shows the two frequency responses reaching almost the same value at the frequency steady state point, that is, about 45,780 MW/Hz for S2-B and 43,955 MW/Hz for S2-C. However, with respect to the 27,126 MW/Hz achieved in S2-A, the improvement is of 18,654 MW/Hz for S2-B and 16,829 MW/Hz for S2-C, values that are significantly larger than the required 7,955 MW/Hz. The reason is the required

frequency response is set to mitigate the lowest frequency due to the under-frequency contingency. As can be seen in Figure 10(a), between 8 to 10 seconds, when the bus frequency reaches its lowest values for the base case S2-A, the frequency responses in Figure 6 improve with about 9,128 MW/Hz for S2-B and 8,173 MW/Hz for S2-C. However, during the frequency recovery period, load is not controlled anymore, as seen in Figure 10(b). Therefore, the load deviation remains the same while the frequency deviation decreases leading to an increase in frequency response.

In the extreme availability scenario, approximately 13M water heaters and other similarly-sized residential loads are scattered through the WECC model the same way as in the other cases, accounting for about 10GW of available frequency controlled load. Comparing the frequency curves in Figure 12(a) at the same representative bus where FRLs are connected, it can be seen that this current scenario presents a behavior opposite to the one in the low availability scenario. The non-supervised frequency control strategy will result in a higher drop in load at minimum frequency (Figure 12(b)), which in turn leads, in this particular setting, to a faster and better frequency recovery compared to the supervised frequency control method.

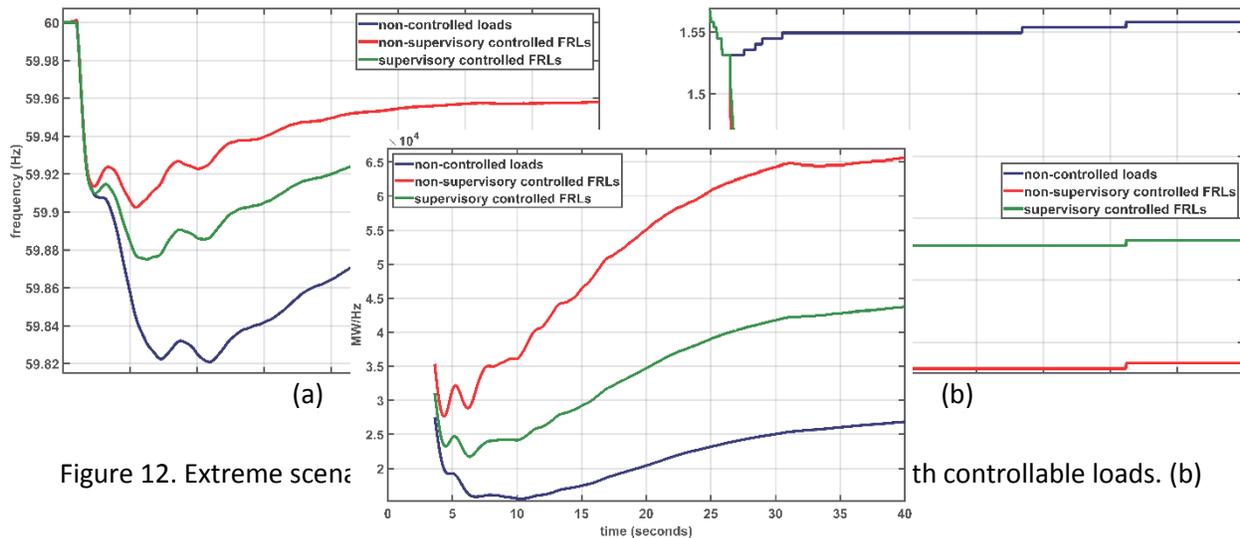


Figure 13. Extreme availability scenario - System frequency response in MW/Hz

Analyzing the frequency responses in Figure 13, the results after the final frequency recovery stage are roughly 65,620 MW/Hz for scenario S2-B and 43,746 MW/Hz for scenario S2-C. This leads to an improvement of 38,758 MW/Hz for S2-B and 16,884 MW/Hz for S2-C, from the 26,862 MW/Hz value for scenario S2-A.

Once again, as in the high availability scenario, the values of the frequency response after frequency recovery are higher than the desired one, for the same reasons as explained previously. Around the time frequency drops to the lowest value, S2-C manages to keep the frequency response to about 8,300 MW/Hz relative to the uncontrolled case, which is close to both the required value and the one registered in the case of high availability. However, at the same time, S2-B shows a deviation of

approximately 19,800 MW/Hz, which is excessively high, meaning the system dropped more load than needed.

So far, the results have been analyzed by comparing the control threshold assignment method influence at different levels of penetration. Figure 14(a) and Figure 14(b) show how the frequency response differs for the same type of frequency threshold determination while the population of frequency-controlled devices increases. If the curtailment thresholds are uniformly distributed between certain fixed limits and independent of any overall system characteristics, such as size of contingency, as the population penetration level increases, more and more of the available ones are subject to being curtailed during an under-frequency event, as seen in Figure 14(a). Though, for the particular case analyzed in this section, this seems to be very advantageous, it could potentially lead to frequency overshoot and/or instability under other circumstances. On the other hand, when a supervisor monitoring the available frequency-controlled loads and the system's needs is to assign the curtailment thresholds to each device based on its power, the results change. As the number of devices increases, there is an upper limit for the amount of loads that is going to eventually be curtailed to meet the requirements. From that point on, the frequency response is not going to change much, as shown in Figure 9(b).

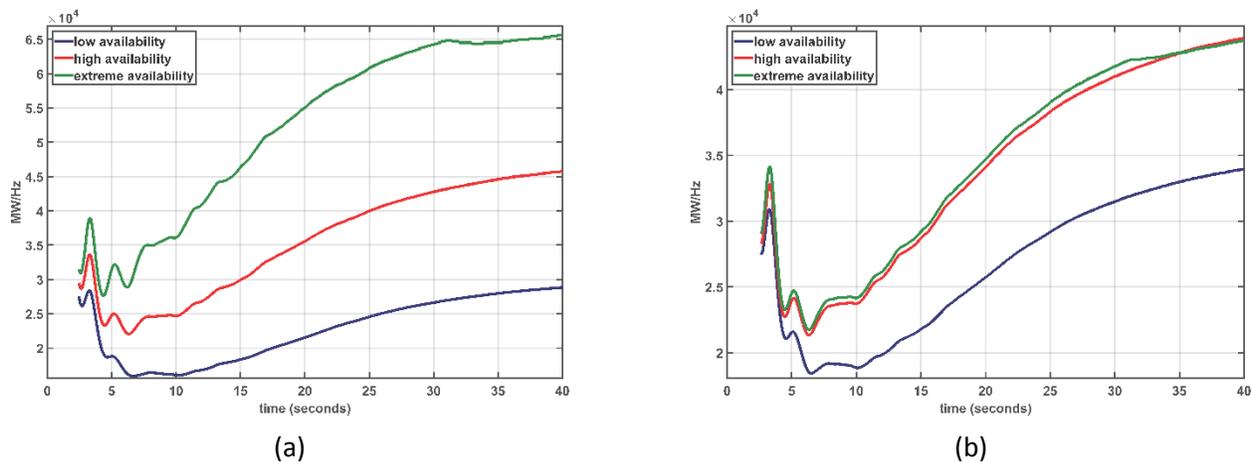
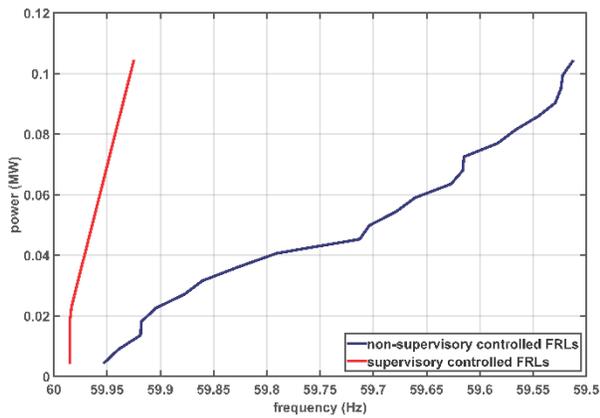


Figure 14. Frequency responses. (a) Decentralized frequency threshold determination. (b) Supervised frequency threshold determination

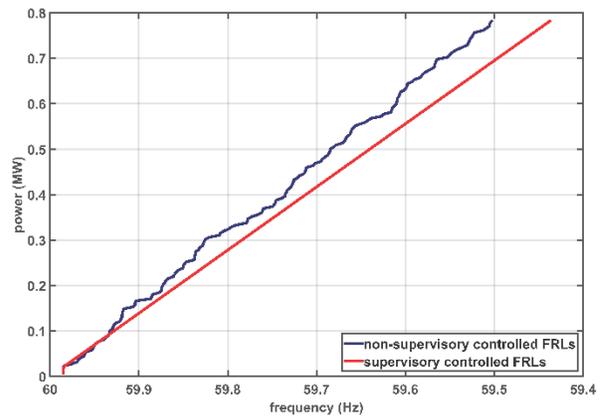
All these results show that there is a strong correlation between the size of the contingency and frequency response to not only the level of FRLs penetration, but also the way the frequency control is set for each available device. For each level of FRLs penetration, that is, low, high and extreme, Figure 15 compares the correlation between frequency thresholds and the amount of power to be dropped if frequency drops below a certain threshold when an under-frequency event occurs.

Figure 15(a) shows that, in the case of low availability, when the thresholds are distributed by a supervisor, rather than prefixed by manufacturers, there will be more participation in frequency regulation from the demand side as compared to the unsupervised case. In this case, some controllers might not get activated due to thresholds lower than the lowest frequency reached during contingency.

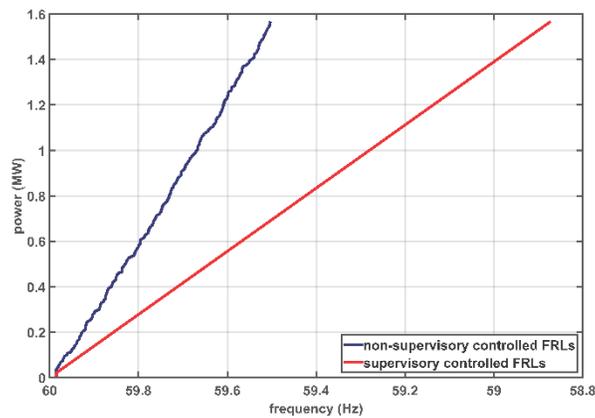
As the penetration level of FRLs increases, the two curves tend to be very similar. Hence, as seen in Figure 15(b), with a highly enough availability of FRLs, the decentralized purely randomized threshold assignment method leads to similar load reduction values as the supervised strategy. However, as shown in Figure 15(c), as the available population increases to extremely high numbers, the slope of the power versus frequency curve keeps increasing for the unsupervised case, possibly leading to too much load drop too early in the transient and risking overshooting and instability. At the same time, the slope of the curve in the supervised case remains almost the same as in the high availability case, and assigns low thresholds to the extra available FRLs. This way, it is ensured that the under-frequency event would not trigger these devices and interfere with other higher-level frequency control actions, such as load



(a)



(b)



(c)

Figure 15. Decentralized vs. supervised threshold determination. (a) Low availability. (b) High availability. (c) Extreme availability. shedding or Remedial Action Scheme (RAS).

4 Investigation of distribution level impacts of using FRLs for providing primary frequency response

In Section 2, a framework for hierarchical supervisory control of FRLs was presented. In Section 3, the proposed framework was tested on the WECC model in PowerWorld. When performing these studies, an aggregate model was used to represent the distribution system at the transmission level. In this section, we will investigate the impacts of the proposed framework on the distribution system. In particular, we will investigate voltage impacts, along with line and transformer overloads as a side effect of turning significant portions of the distribution load on or off to restore the system frequency back to nominal after a contingency.

The section is divided into four sub-sections. In Section 4.1, the specific distribution test system model used to perform the analysis is discussed, followed by the base case used in the studies in Section 4.2. Section 4.3 will show the effect on voltage and line overloads of using water heaters for primary frequency control. Possible voltage protection algorithms to mitigate the distribution voltage impacts are discussed in Section 4.4.

4.1 IEEE 8500-Node Test System

The hierarchical supervisory control of FRLCs proposed in Section 2 is implemented in GridLAB-D. GridLAB-D is a multi-agent simulation and modeling environment for engineered systems, with particular emphasis on power and energy systems. In this work, GridLAB-D is used to simulate the IEEE 8500-Node Test System. This system was created by the IEEE Distribution Analysis Subcommittee's Test Feeder Working Group [18]. It was designed to provide a realistic benchmark for the analysis community; an actual distribution circuit with significant load and voltage imbalances and multiple inline voltage control devices [19], [20]. It contains three controllable capacitors and three line regulators, with a peak load of approximately 11.9 MW. In most of the IEEE test systems, static power flows are used with simple load models; however, in order to demonstrate the effectiveness of the FRLCs in providing primary frequency control, a more dynamic model is needed. The simple loads are replaced with physics-based load models [21], including a state-based model of the residential HVAC system and water heater [22]. The water heater models are driven by water usage schedules, randomized for each device to represent standard water draws (e.g., a shower versus a dishwasher cycle), and following aggregate load shapes determined from ELCAP data. The number of residential units placed on the circuit is determined by calibrating the load models to the peak load. The entire load allocation methodology is described in more detail in [23]. The circuit considered in the studies includes 1,977 residential homes with electric water heaters.

Along with detailed models of the water heater behavior, we also implemented the proposed framework in Section 2.1 as a controller that is attached to each water heater, able to alter the state operation of the water heater in the presence of an emergency frequency event. This controller further communicates with a supervisor object, as described in Section 2.3, implemented in order to coordinate the distribution of response frequencies for the individual devices. In addition to these changes, timer-

based voltage control settings had to be included. Standard utility practice guidelines were used to create a model that is operated in line with common utility practices:

- average primary voltage was targeted at approximately 126 Volts equivalent, assuming a 3-4 Volt drop from the primary to the customer meter
- control delays cascaded down the circuit using standard timing for regulator operations
- circuit-level power factor was maintained between 0.9 lagging and 1.0
- controller set points were set to minimize capacitor and regulator interactions (capacitors primarily targeted reactive power compensation, not voltage)

The regulator and capacitor setpoints are shown in Table 1.

Table 1. Regulator and capacitor settings (modified from original system)

Regulator and capacitor settings use in 8500 node system				
	Regulator 1	Regulator 2	Regulator 3	Regulator 4
Location	Substation	North branch; Furthest from substation	North branch; between Reg1 and Reg2	South branch
Voltage Setpoint	125	124.67	124.67	125
Time Delay (s)	60	120	75	90
Capacitor settings				
	Capacitor 1	Capacitor 2	Capacitor 3	Capacitor 4
Location	North Branch; Between Reg2 and Reg3	Downstream from substation; Before branching	Substation	South branch
kVAR High	475	425	450	FIXED
kVAR Low	-350	-350	-350	FIXED
Voltage High	130	128	128	FIXED
Voltage Low	114	114	114	FIXED
Time Delay (s)	480	300	180	FIXED

*See Figure 10 for a map of the circuit.

The resultant circuit is representative of standard utility distribution network operations, maintaining voltage within acceptable parameters and operating voltage control devices a reasonable

number of times per day. This circuit was chosen primarily due to the voltage control issues which existed prior to deploying new distributed control technologies. The large number of voltage control devices is not typical to all distribution utilities, although it is not uncommon. It is expected that this circuit represents an extreme (although realistic) case.

4.2 Case 0 - Base Scenario

In this section, we will discuss settings and results for the distribution system described in Section 4.1. This baseline will be used to compare other results with. For this baseline, we will not be imposing any frequency contingency. We simulated a four-hour window in a mild, but warm September day using TMY2-based Yakima, WA weather [24]. We simulated between 12:00 – 16:00, with most results showing only 14:00 – 16:00. This shortened window was used due to the agent-based nature of the simulation requiring time to settle to steady state, which occurs in a little over an hour from the start of the simulation. As mentioned previously, the system has a peak load of approximately 11.9 MW. Since we are simulating a milder day, we are not at peak load, as seen in Figure 16. The total load for the system is measured at roughly 7.1 MW, and of that, roughly 2 MW is from electric water heaters.

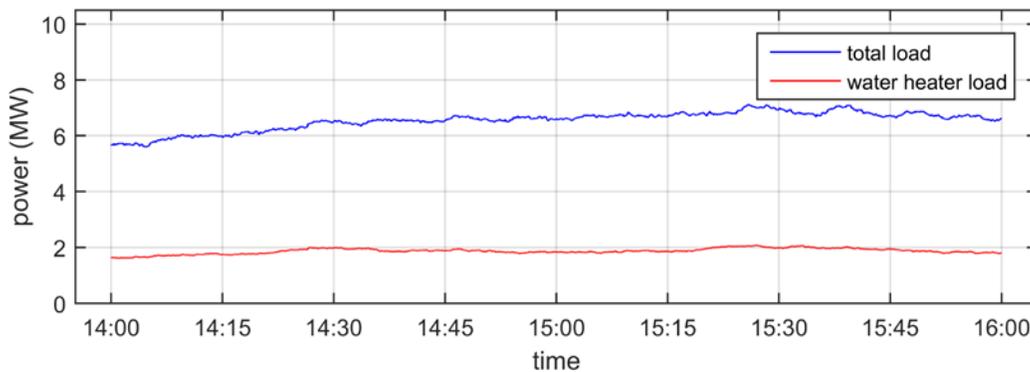


Figure 16 Load profile for the IEEE 8500 node system.

As mentioned earlier, there is a total of 1977 households simulated, each of them having an electric water heater. Not all distribution circuits have 100% penetration of electric water heaters, but this is used to show the extreme event, i.e., the most load that can be synchronized. Figure 17 shows the number of water heaters on between the hours of 14:00 and 16:00. From Figure 17, we see that around 20% of the total population of water heater are on during this period.

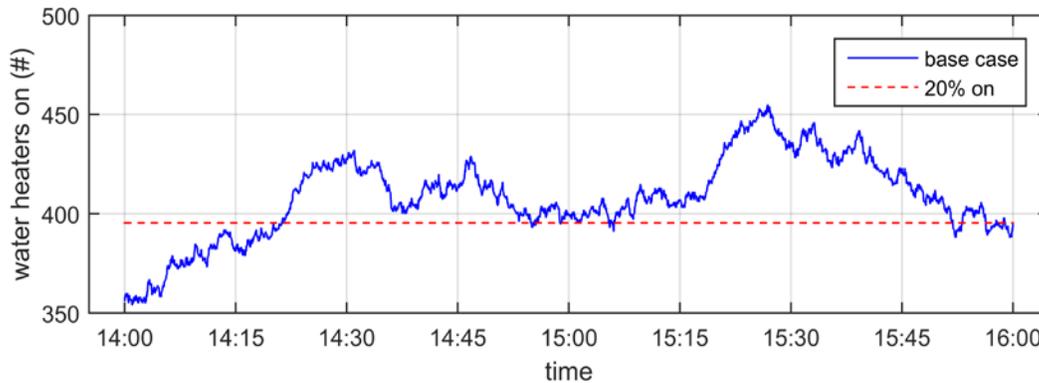


Figure 17 Number of water heaters that are on.

4.2.1 Distribution System Metrics

Distribution utilities have the responsibility for maintaining service and reliability within the distribution system. When resources are being controlled, and particularly synchronized, outside of the control of the local reliability coordinator, there is the potential for causing unforeseen issues on the distribution system. To quantify these impacts, a series of metrics related to distribution operations were created. These should not be thought of as “pass/fail”, but rather a quantification of the impacts. Specific distribution utilities would need to determine whether the impacts are acceptable (or not), considering the infrequent nature of these events.

The first two metrics are related to maintaining voltage within acceptable bounds at the point of customer interconnect. ANSI C84.1 requires that the service voltage (at the customer meter) be maintained within a 5% limit of the 120 V nominal (114 V – 126 V). This is called Range A. Range B is designated for emergency situations, and allows for a slight extension of the range (110 V – 127 V) for a short period of time. While not specifically designated by ANSI, instantaneous voltages should not exceed a 10% limit at any given time (108 V – 132 V), as this can cause equipment damage.

In the simulation of the 8500 node test system, voltage is measured at every customer connection at every second. At each connection point, if voltage exceeded Range A requirements, either high or low for five minutes, it was considered a “continuous violation”. If the voltage exceeded 10% of nominal, high or low, this was considered an “instantaneous violation”. Again, these metrics are not used to indicate that the circuit “failed”, but rather an indication of the level of impact.

In Figure 18, the voltage profile for the base case is shown, as measured at the customer meter. This is a snapshot of the voltage at time 14:31. In addition, voltage control equipment, including voltage regulators (red dots) and capacitors (blue dots), are shown. Rough approximations of voltage control regions and feeder topology are also shown in the figure – in later figures, for clarity, the control regions and topology will not be shown. Note, that the circuit is being operated at approximately 123 V (1.025 puV) at the customer meter, but that there is also a wide spread of voltages ranging from 0.96 to 1.05 puV. This is a fairly common utility practice; utilities will operate at the high end of the ANSI band to ensure that no voltage drops below the safe threshold. It should also be noted that this circuit was chosen because it has significant voltage control issues prior to deploying distributed resources. It is

likely that any changes to the behavior of this circuit will significantly impact the voltage control and may not be representative of all circuits.

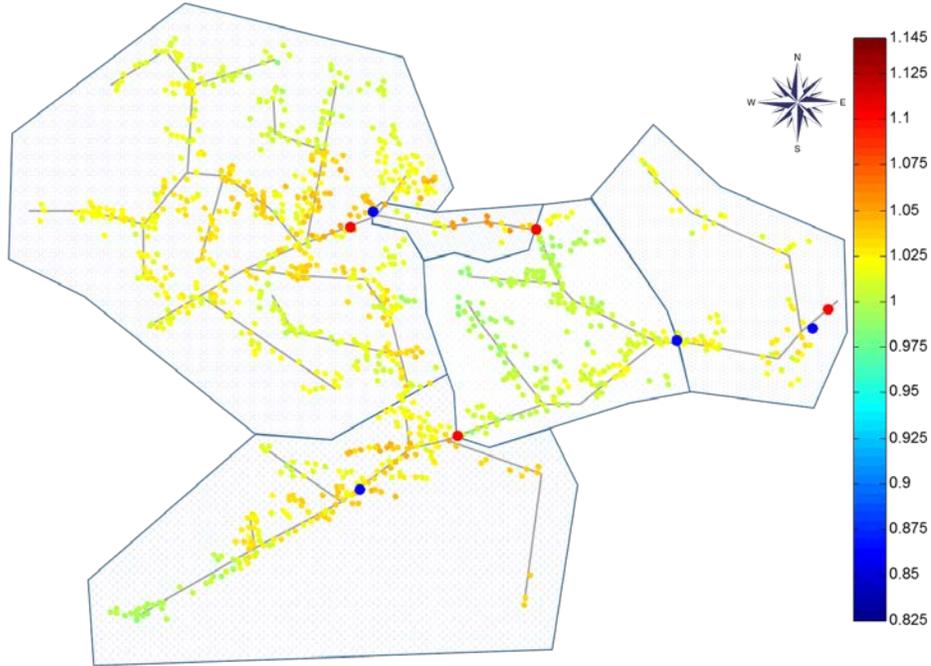


Figure 18 Voltage profile for the base case taken at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks.

Table 2, shows that the base case indicates zero instantaneous voltage violations and four meters with at least one continuous high voltage violation.

Table 2 Voltage violations for base case

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	4	0	0	0

The third metric is related to equipment overloads, particularly excessive over-current through overhead lines, underground cables, and service transformers. Current flow through each of these devices is measured every second in the simulation. If the current flow of the underground cables or overhead conductors exceeds the designated rating (i.e., the over-current limit), a violation is flagged. For the transformers, if current exceeds 200% of rated, then a violation is flagged. Transformers can operate in over-current conditions, as long as sufficient time is given for the device and oil to cool off;

200% overload is used as an arbitrary limit to indicate excessive wear-and-tear. In all cases presented, the following transformer and line violations stayed the same and can be attributed to improper adjustment of the baseline. A summary of these overloads is presented in Table 3.

Table 3 Transformer and line overloads

Device	Violation Count
Transformer T5338989A	17
Transformer T226192762B	1447
Transformer T5223658A	171
Transformer T226196642A	216
Line Tpx227944551B0	46

The fourth metric is related to the number of tap changes for voltage regulators and the number of switches for capacitor banks, particularly excessive movement. The state of the voltage regulators and capacitor banks is measured every second in the simulation and the number of state changes is tallied. Results from the base case can be seen in Table 4 for the time windows between 14:00 and 16:00. Note, that this equates to approximately 50 tap changes per day, which is a little on the high side. However, considering the voltage control issues of this circuit, it is not unreasonable. The number of capacitor state changes is low, since these devices are primarily used for reactive power control, while voltage control is a secondary consideration.

Table 4 Number of voltage regulator tap changes and capacitor bank switching

Device	Phase A	Phase B	Phase C
Regulator VREG2	5	6	5
Regulator VREG3	8	9	6
Regulator VREG4	8	7	9
Capacitor CAPBANK0	1	1	0
Capacitor CAPBANK1	0	0	0
Capacitor CAPBANK2	0	0	0

In each of the following scenarios, these metrics will be used to describe the impacts to distribution system operations. Only those that had significant changes from the base case will be discussed.

4.3 Case 1 – Supervisory Control of Distributed FRLs

In this case, we will run with the same settings as base case, except device-level controllers are added to each water heater and a supervisory controller of the distributed FRLs is activated. As described in Section 2, this means that we will have all water heater monitoring frequency and responding accordingly in the event of an emergency frequency event. The first frequency event we will be analyzing is depicted in Figure 19, and corresponds to tripping two large generators in the south of the WECC system. This contingency is the same one used in Section 3.

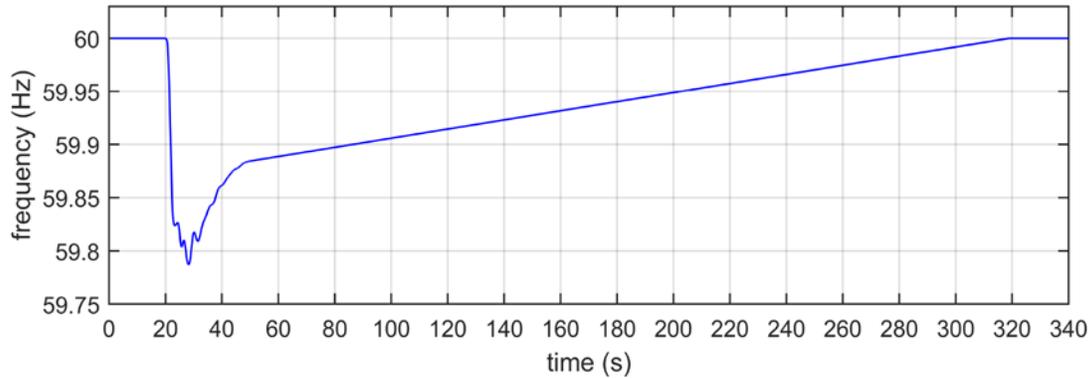


Figure 19 Under-frequency event from tripping two generator in the south of WECC (~2.7 GW)

The total amount of generation tripped during this event is roughly 2.7 GW. The frequency signal is constructed based on 30 seconds of data from a PowerWorld simulation. Due to the PowerWorld model lacking AGC control and long term dynamics, we are assuming that AGC control will take over after 30 seconds and linearly bring the frequency back to nominal within 5 minutes.

With the frequency signal described above, we are able to adjust the supervisory droop to get the desired response from the population of water heaters. The supervisor adjusts the droop-like curve every 15 minutes and we assume, for all simulations, that the supervisor is calculated right before the contingency. We are adjusting the droop to get 2 MW of response with this contingency, which corresponds to a droop setting for the supervisor of 10%. In Figure 20, the response curve for Case 1 with the under frequency event is reported. In this plot we have both the expected curve based on the supervisor settings and the observed one. For this case we see that we get the expected response.

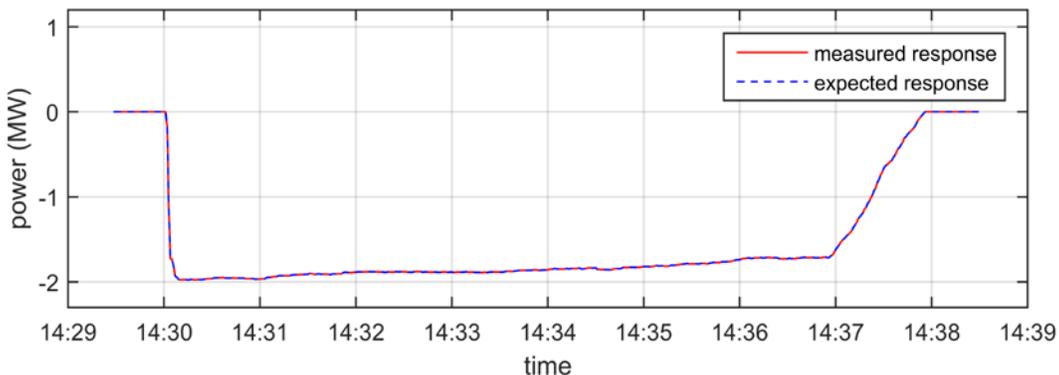


Figure 20 Response in MW for Case 1 with under-frequency event.

In Figure 21, the effect on the voltage profile when the FRLCs react to the frequency event is shown. The snapshot of the voltage is taken one minute after the contingency (same time as the base case snapshot). It is clear that the voltage at some customer meters are outside the emergency range and the overall profile is very different than the one in Case 0. Devices marked with '+' are the ones changing state due to the contingency. In this case, as expected these device are distributed evenly throughout the system.

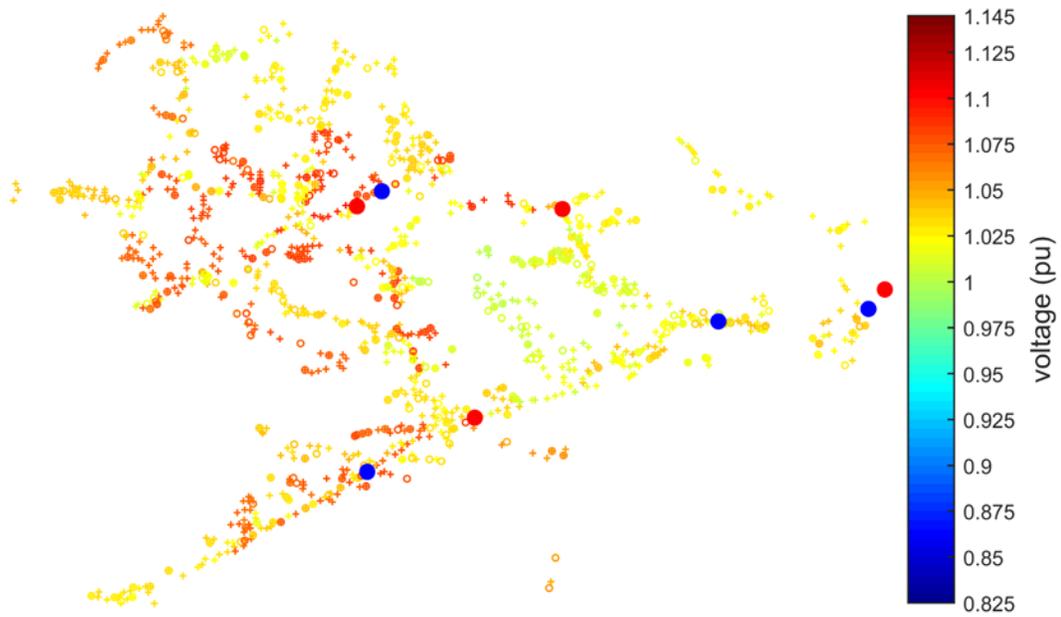


Figure 21 Voltage profile for Case 1 with under frequency event taken at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.

From Table 5, we see that we have 1648 instantaneous high-voltage violations along with 117 continuous high-voltage violations. This clearly shows that providing frequency response can have a significant impact on the distribution system voltage operations. For Case 1, we did not see a significant increase in tap changes for both the regulators and capacitor banks, as the event is short lived, or any additional line or transformer overloads.

Table 5 Violation data for Case 1 during under-frequency event.

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	117	0	1648	0

The second frequency event we will be analyzing is depicted in Figure 22 and corresponds to tripping one large load in the southern portion of the WECC system. The total amount of load tripped is roughly 0.9 GW. The frequency signal is constructed based on 30 seconds of data from a PowerWorld

simulation. Again due to the PowerWorld model lacking AGC control and long term dynamics, we are assuming that AGC control will take over after 30 seconds and linearly bring back the frequency to nominal within 5 minutes.

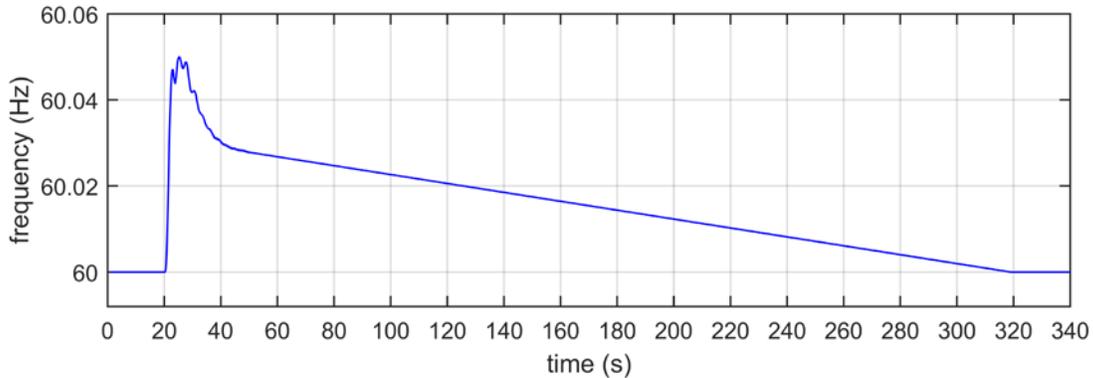


Figure 22 Over-frequency event from tripping large load in the south of WECC (~0.9 GW)

As in the case with the under-frequency event, we will adjust the supervisory droop to get the desired response from the population of water heaters. We are again adjusting the droop to get 2 MW of response with this new contingency, which corresponds to a droop setting for the supervisor at 2.5%. In Figure 23, the response for Case 1 with the over frequency event is reported. In this plot we have both the expected curve based on the supervisor settings and the observed one. For this case, we see that we get the expected response.

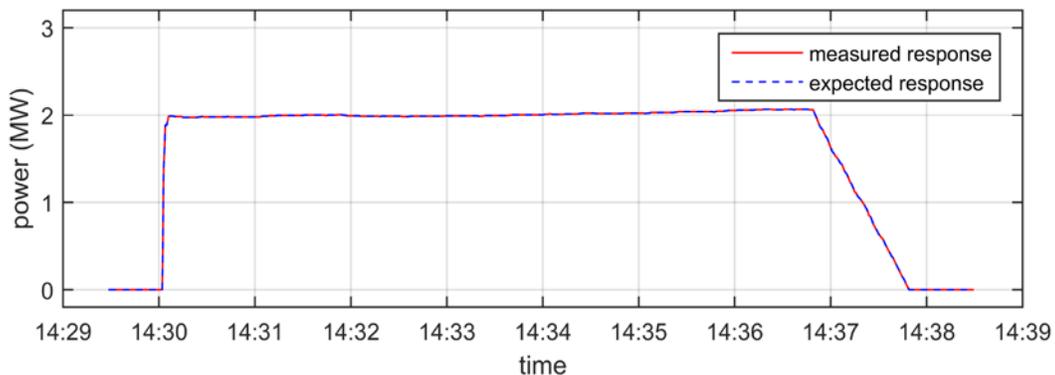


Figure 23 Response in MW for Case 1 with over frequency event.

In Figure 24 the effect on the voltage profile when the FRLCs react to the over-frequency event is shown. The snapshot of the voltage is taken one minute after the contingency. Due to the high baseline voltage level, we have more room before we hit the lower ANSI bounds (as opposed to the high-frequency event that hit the upper ANSI bound). This results in fewer violations.

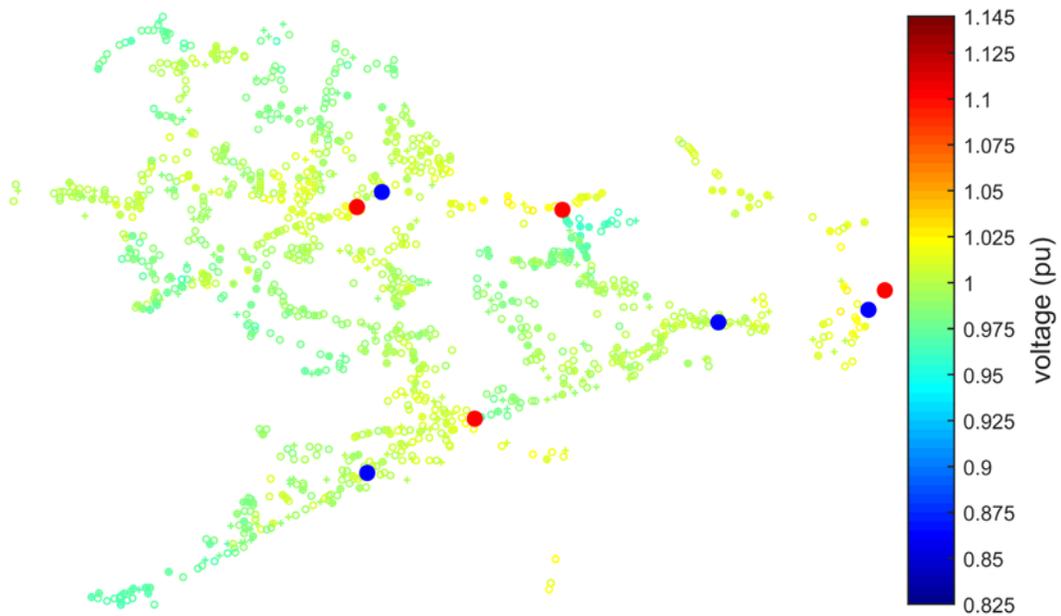


Figure 24 Voltage profile for Case 1 with over-frequency event taken at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.

This is further supported by Table 6, where we see that we do not have any instantaneous voltage violations. In this case, we have 267 continuous high-voltage violations. However, even though we do not have any emergency violations, using load to respond to frequency clearly has an impact on the distribution operations.

Table 6 Violation data for Case 1 during over-frequency event.

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	267	0	0	0

Some of these violations may be due to the increase in regulator tap changes and capacitor bank switching operations experienced during this event. In Table 7, an overview of the additional state changes can be found (positive value indicates increase operations). In the case of an over-frequency event, we have additional state changes from voltage regulators and capacitor banks. We especially see that we have movement of capacitor banks that could potentially create unwanted oscillations and interactions with the voltage regulators. This level of increase would be of significant concern if it were occurring all of the time, as it could impact the lifetime of these devices. However, considering this is a

two-hour window on a single day during a very infrequent event, it likely has little impact on the lifetime of the voltage control devices.

Table 7 Capacitor switches and regulator tap changes for Case 1 during the over-frequency event.

Device	Phase A	Phase B	Phase C
Regulator VREG2	7	16	-2
Regulator VREG3	11	24	-2
Regulator VREG4	8	20	-2
Capacitor CAPBANK0	0	4	0
Capacitor CAPBANK1	0	4	0
Capacitor CAPBANK2	0	0	0

4.4 Case 2 - Supervisor Control of Distributed FRLs Considering Voltage Impacts of the distribution system

In Case 1, we presented initial distribution system impacts of providing primary frequency response. In this subsection, we will describe initial measures taken to mitigate these effects. As the primary “failure” mechanism, the additional control algorithms developed will focus on decreasing voltage impacts and are explained in the following section.

The first algorithm developed is a voltage lockout of participating devices. This is a local control algorithm that is not communicating with the supervisor. In this algorithm, each device will measure the voltage at the outlet at every control period and compare it to predetermined thresholds. If the measurement is outside the defined thresholds, the control output of the FRL will be overridden and the device will be returned to normal operation. Furthermore, a wait period is also implemented as to mitigate excessive switching.

The second algorithm developed sorts the devices as they are placed in the supervisor, rather than treating every device equally. Two different approaches to sorting the incoming bids according to voltage were investigated. Approach A sorted bids according to ascending absolute voltage deviation from nominal. This will ensure that devices closest to nominal voltage will receive frequency thresholds closest to nominal frequency, and that those devices nearest the ends of the ANSI band will be the last to act. The idea behind this being that the parts of the feeder with voltage problems will be used last in the event of an emergency frequency event. Approach B, will divide the populations into two bins, devices that are *on* will be sorted such that the devices with the highest voltage magnitude will receive frequencies closest to nominal frequency, and devices that are *off* will be sorted such that the devices with the lowest voltage magnitude will receive frequencies closest to nominal frequency. The idea behind this being that the parts of the feeder with the highest voltage level will be used first in the case of an over-frequency event and devices with the lowest voltage level will be used first in the case of an under-frequency event.

4.4.1 Case 2-1 Supervisor Control of Distributed FRLs with Voltage Lockout

In Case 2-1, we used the same settings as Case 1, with the exception of the activation of the local voltage lockout method described previously. We set the voltage lockout at 4%, meaning that if voltage deviated more than 4% from nominal, we consider the device to be held in a locked out state for a period of 60 seconds. For the simulation, we used the same droop setting as in Case 1, where we expect 2 MW of response with the under-frequency contingency. In Figure 25, the response curve for the Case 2-1 under-frequency event is shown. In this plot, we have both the expected curve based on the supervisor settings and the observed one. From the figure, it is clear that having devices entering a voltage lockout state will impact the response of the devices. In this case, we see a reduction in the initial response of roughly 25%.

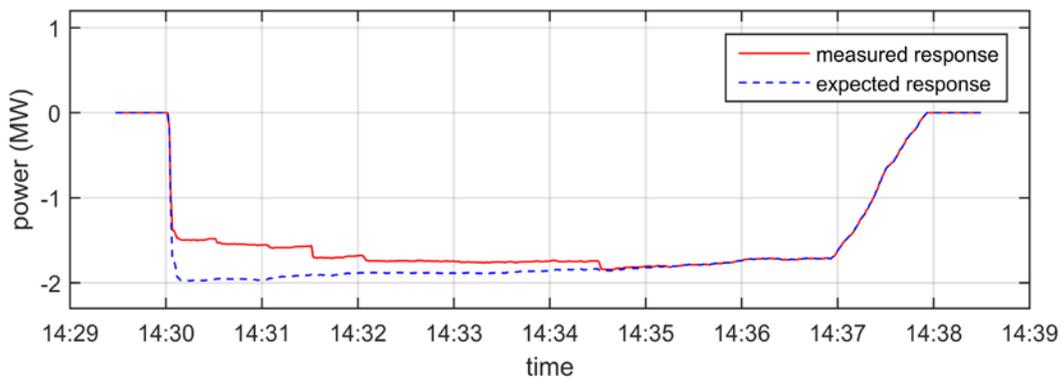


Figure 25 Response in MW for Case 2-1 with the under-frequency event.

As seen in the response curve we have a significant portion of devices in voltage lockout. Due to this we will expect to see a difference in the voltage profile for the system. In Figure 26, this profile is shown and we see that compared to Case 1, the voltage levels across the system are decreased from the case without the voltage lockout enabled.

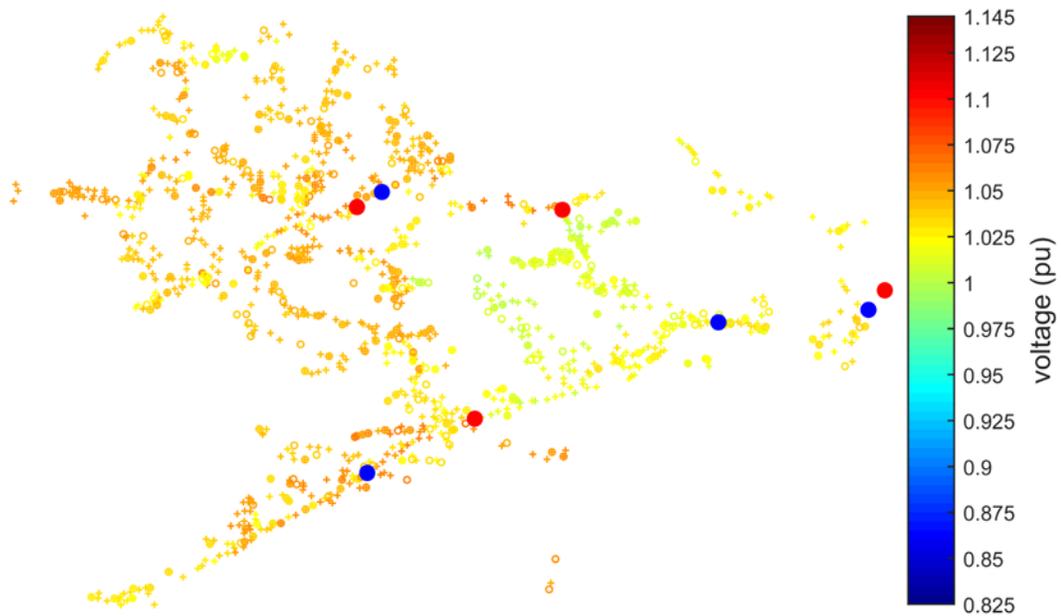


Figure 26 Voltage profile for Case 2-1 with the under-frequency event taken at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.

In order to better quantify the effects over a period of time, Table 8 summarizes the voltage violations. We see that, compared to Case 1, we have eliminated all of the instantaneous voltage violations. We also see a decrease in the continuous violations from 117 to 49. This confirms that for this specific case, voltage lockout is helping to stabilize the distribution system voltages.

Table 8 Violation data for Case 2-1 during the under-frequency event.

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	49	0	0	0

With the same settings, we ran Case 2-1 using the over-frequency event reported in Case 1. Figure 27, shows the response for Case 2-1 over-frequency event. In this plot, we have both the expected curve based on the supervisor settings and the observed one. From the figure, it is clear that having devices in voltage lockout is again impacting the response. In this case, we see both a difference in the initial response and during the event (at time 14:34.5) with a difference greater than 25% from the expected response.

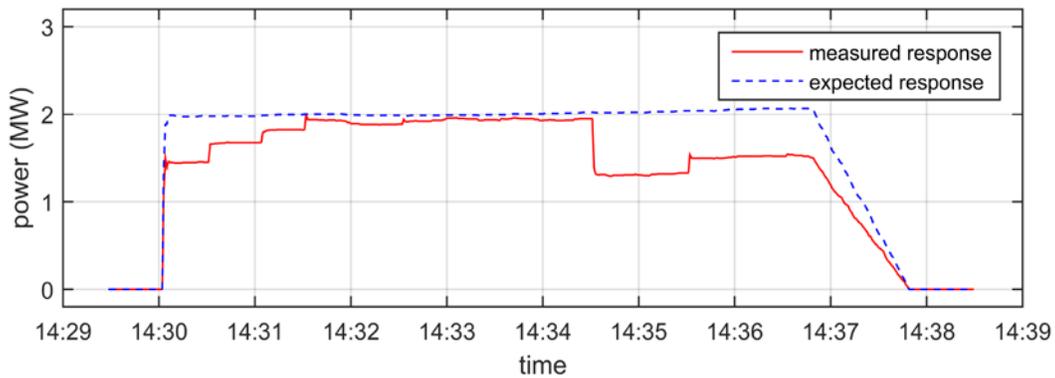


Figure 27 Response in MW for Case 2-1 with the over-frequency event.

This tells us that a significant portion of devices are in voltage lockout during different times of the event. Due to this, we will expect to see a difference in the voltage profile for the system. In Figure 28, this profile is shown and we see that, compared to Case 1, there is a slight decrease in the overall voltage level.

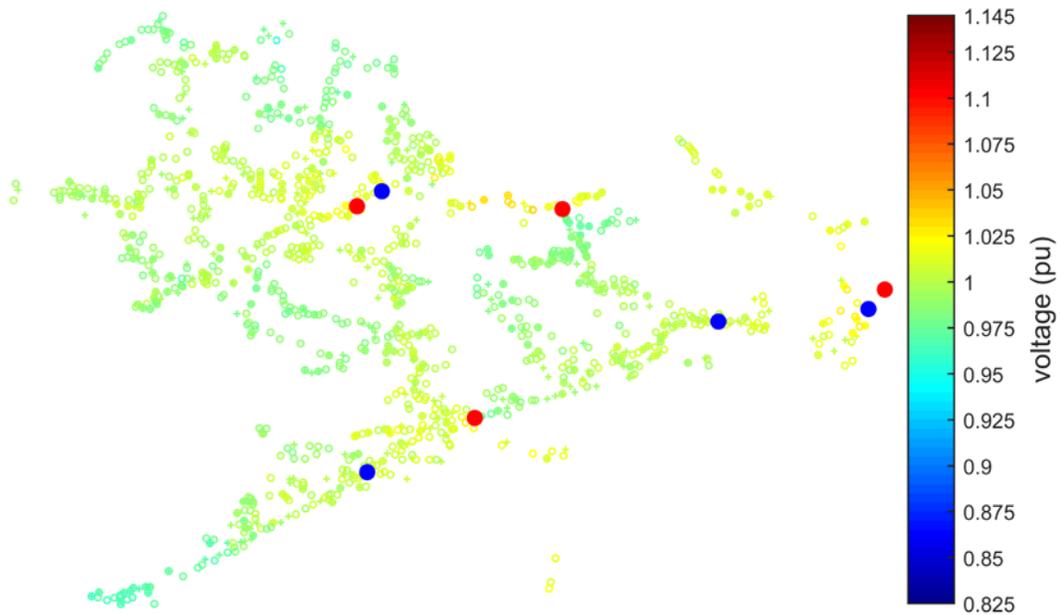


Figure 28 Voltage profile for Case 2-1 with the over-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.

However, looking at the violation summary tells a different story about the effectiveness of the voltage lockout. In Table 9, we see that, compared to Case 1, we have an increase in instantaneous voltage violations from 0 to 1459. We also see an increase in the continuous violations from 267 to 322.

Table 9 Violation data for Case 2-1 during the over-frequency event.

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	322	0	1459	0

Looking a little further reveals the cause of the increase in violations is due to a significant increase in voltage regulator taps and capacitor bank switching. A summary of the voltage control state changes are reported in Table 10.

Table 10 Regulator tap changes and capacitor switches for Case 2-1 during the over-frequency event.

Device	Phase A	Phase B	Phase C
Regulator VREG2	8	16	0
Regulator VREG3	13	22	0
Regulator VREG4	6	18	2
Capacitor CAPBANK0	0	4	0
Capacitor CAPBANK1	0	4	0
Capacitor CAPBANK2	0	0	0

In Figure 29, the evolution of states for capacitor bank 1 (phase B) and voltage regulator 2 (phase A, B, and C) are reported. From Figure 29, it is clear that the frequency event triggers a slow-dynamic oscillation between the capacitor and voltage regulator control. This oscillation significantly increases the number of tap changes and capacitor switching operations. The oscillations persist for almost one hour, until the system can settle back into equilibrium. This oscillation is the reason for the increase in voltage violations. It should be noted that this number of operations is unlikely to cause voltage control device lifetime degradation, due to the infrequency of these events. However, this type of rapid “chattering” of voltage control devices could lead to enough voltage deviations that customers may complain to the utility about “flicker” due to poor power quality [25]. Again, due to the infrequency of these type of events, this may be of little concern to the distribution utility, but may require customer education.

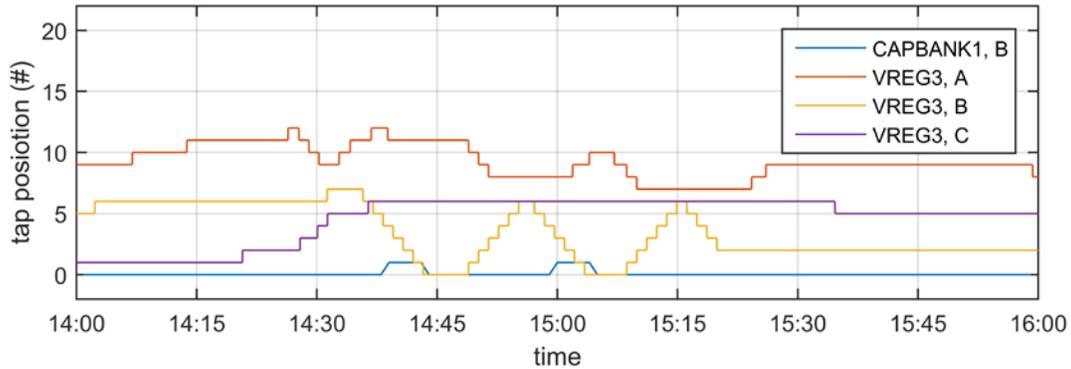


Figure 29 Tap changes for VREG3 (phase A, B, and C) along with switching operations for CAPBANK1 (phase B) during the over-frequency event with voltage lockout enabled.

4.4.2 Case 2-2 Supervisor Control of Distributed FRLs with Voltage Sorting Approach A

In Case 2-2, we will use the same settings as Case 1, with the exception of activating the voltage sorting Approach A in the supervisor. For the simulation, we will be using the same droop setting as in Case 1 where we expect 2 MW of response with the under-frequency contingency. Since we do not impose any voltage lockout in this case, we see the same power (MW) response as in Case 1.

In Figure 30, we show the voltage profile for the case with voltage sorting activated. Again, the snapshot of the voltage is taken one minute after the contingency. It is clear that the voltage at some customer meters are outside the emergency range. When compared to Case 1, the voltage profile is very similar.

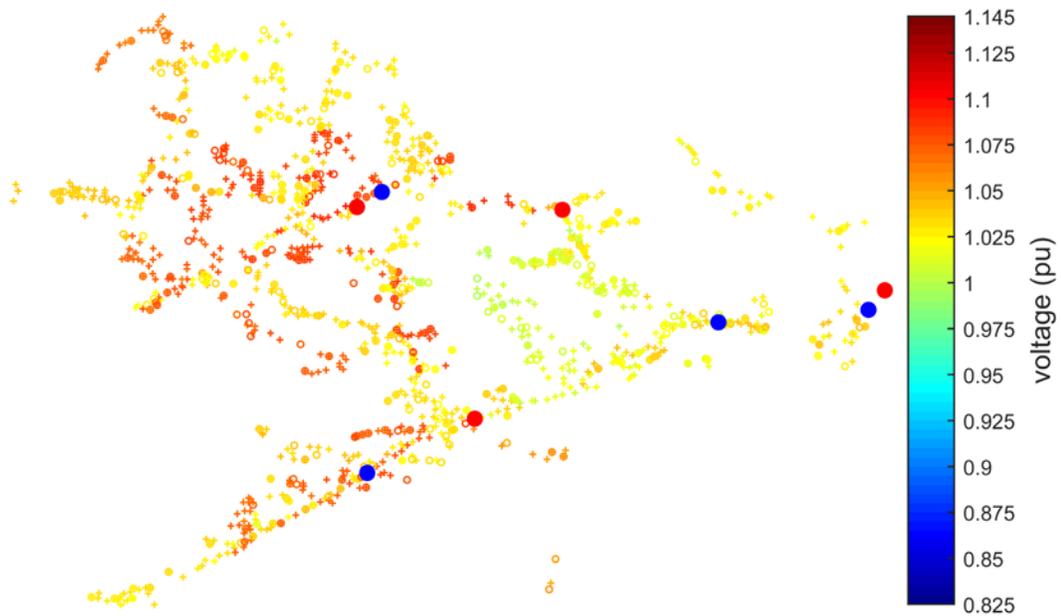


Figure 30 Voltage profile for Case 2-2 with the under-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.

This is further confirmed in Table 11, where the voltage violations are reported. In this table, it is clear that using the voltage sorting Approach A does decrease the violations, slightly. One reason for the large number of violations is that for this specific under-frequency contingency, most of the devices that are *on* are turned *off* to achieve the 2 MW response requirement. When only a few devices are left *on*, the sorting of voltage is not as effective since most units, no matter their frequency threshold, will be activated.

Table 11 Violation data for Case 2-2 during the under-frequency event.

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	103	0	1605	0

In Case 2-2 with the over-frequency contingency, we will be using the same droop setting as in Case 1. In this case we will again be asking for 2 MW of response. Since we do not impose any voltage lockout in this case we see the same power (MW) response as in Case 1.

In Figure 31, the voltage profile for Case 2-2 with the over-frequency event is reported. In this, it is clear that using voltage sorting Approach A in the supervisor does not have the desired effect during this contingency. Compared to Case 1 we see a voltage profile with a generally higher voltage. One thing to note in this plot is the tendency of devices to group and switch states simultaneously during the contingency. These devices are marked with the symbol '+'. The reason for this lies in the nature of the specific voltage behavior of this system. By sorting voltage according to ascending absolute deviation from nominal, we are calling upon devices that prior to the contingency are close to nominal voltage. Due to the voltage control regions (shown earlier), this means that the devices are grouped topologically and have a greater impact on local circuit voltage when synchronized.

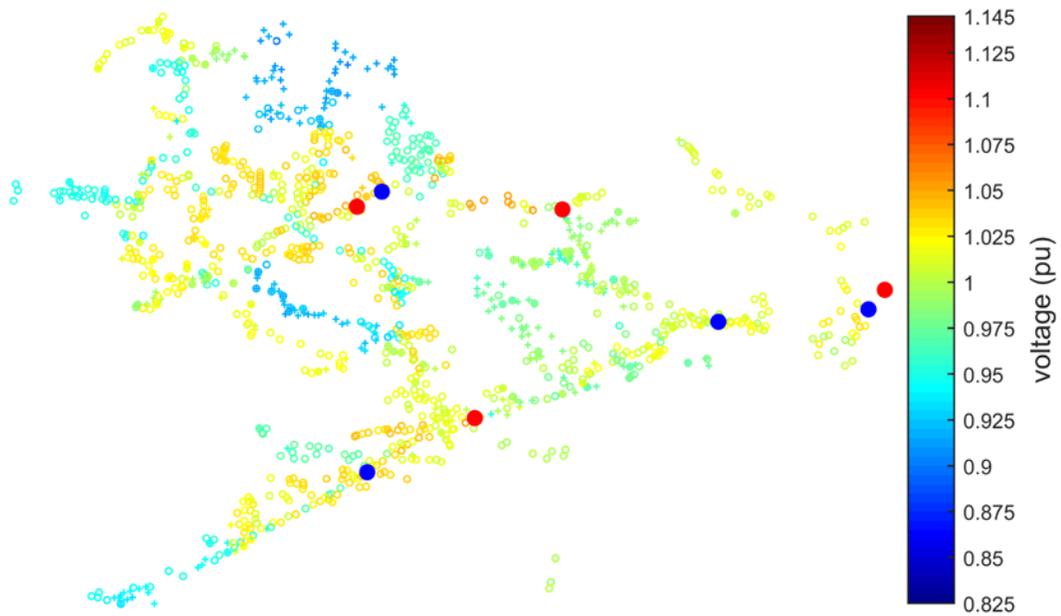


Figure 31 Voltage profile for Case 2-2 with the over-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event. The above conclusion is further supported by the violation summary in Table 12.

Table 12 Violation data for Case 2-2 during the over-frequency event.

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	478	0	475	39

4.4.3 Case 2-3 Supervisor Control of Distributed FRLs with Voltage Sorting Approach B

In Case 2-3, we will use the same settings as Case 1, with the exception of activating the voltage sorting Approach B in the supervisor. For the simulation, we will be using the same droop setting as in Case 1 where we expect 2 MW of response with the under-frequency contingency. Since we do not impose any voltage lockout in this case, we see the same power (MW) response as in Case 1.

In Figure 32, we show the voltage profile for the case with voltage sorting activated. Again, the snapshot of the voltage is taken one minute after the contingency. It is clear that the voltage at some customer meters are outside the emergency range and that compared to Case 1, the voltage profile is very similar.

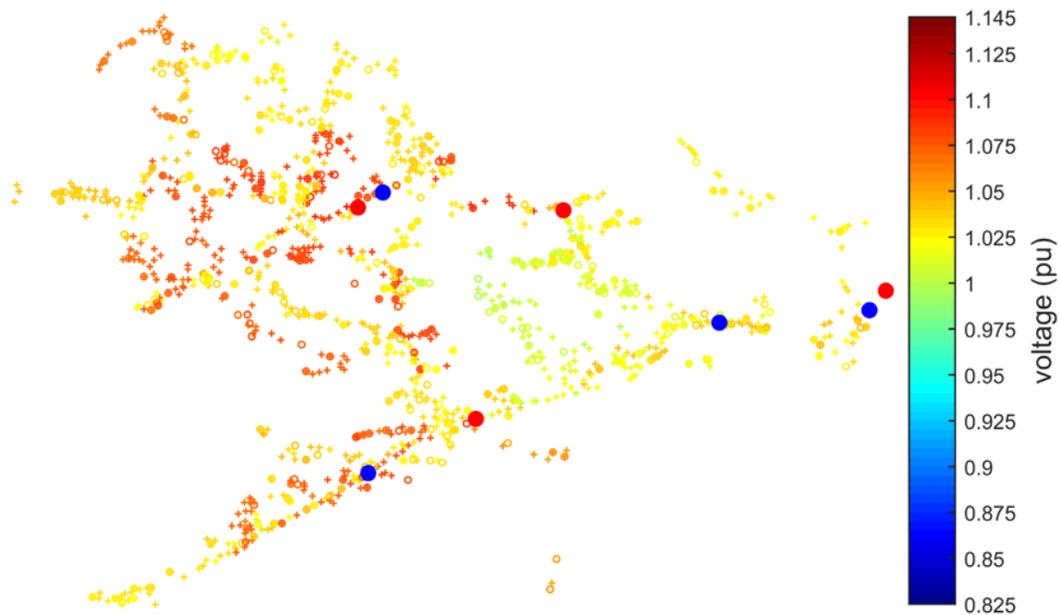


Figure 32 Voltage profile for Case 2-3 with the under-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.

This is further confirmed in Table 13, where the voltage violations are reported. In this table, it is clear that using the voltage sorting method does decrease the violations slightly. One reason for the large amount of violations still present is that for this specific under-frequency contingency, most of the devices that are *on* are turned *off* to achieve the 2 MW response requirement. When only a few devices are left *on*, the sorting of voltage is not as effective since most units, no matter their frequency threshold, will be activated.

Table 13 Violation data for Case 2-3 during the over-frequency event.

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	103	0	1605	0

In Case 2-3 when using the over-frequency contingency, we will be using the same droop setting as in Case 1. In this case, we will again be asking for 2 MW of response. Since we do not impose any voltage lockout in this case, we see the same power (MW) response as in Case 1.

In Figure 33, the voltage profile for Case 2-3 with the over-frequency event is shown. In this, it is clear that performing voltage sorting in the supervisor does not have the desired effect during this contingency. Compared to Case 1, we see a voltage profile with a greater discrepancy from nominal. In some areas, the voltage is very high and in others very low. One thing to note is the tendency of devices to group topologically and switch states simultaneously during the contingency. These devices are marked with the symbol '+'. The reason for this lies in the voltage control regions which tends to cause the devices to group by voltage, and therefore, topologically. By sorting the devices via voltage according to Approach B, we are calling upon devices that prior to the contingency have high voltages and are lumped together on certain portions of the circuit. This synchronizes loads on certain portions of the circuit causing drastic voltage deviations, and leads to "hunting" by the voltage control devices as they try to bring the circuit voltage back to nominal.

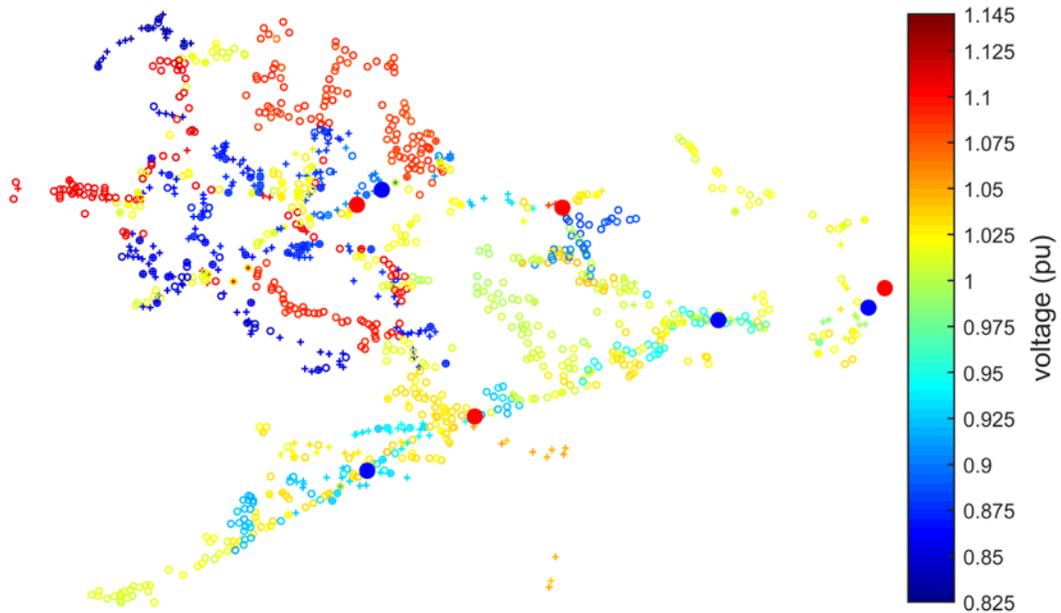


Figure 33 Voltage profile for Case 2-3 with the over-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.

The above conclusion is further supported by the violation summary in Table 14. Here we see the number of violations have significantly increased in both metrics.

Table 14 Violation data for Case 2-3 during the over-frequency event.

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	653	581	67,672	139,219

4.4.4 Case 2-4 Supervisor Control of Distributed FRLs with Voltage Lockout and Voltage Sorting Approach A

In Case 2-4, we will use the same settings as Case 1, with the exception of the activation of both local voltage lockout and voltage sorting Approach A in the supervisor. In the following, we will have the voltage lockout set at 4% and a lockout of 60 seconds. For the under-frequency contingency, we used the same droop setting as in Case 1, where we expect 2 MW of response. In Figure 34, the response

curve for Case 2-4 with the under-frequency event is reported. In this plot, we have both the expected curve based on the supervisor settings and the observed curve. From the figure, it is clear that having devices in voltage lockout will impact the response from the devices. In this case, we see a difference in initial response of roughly 25% less.

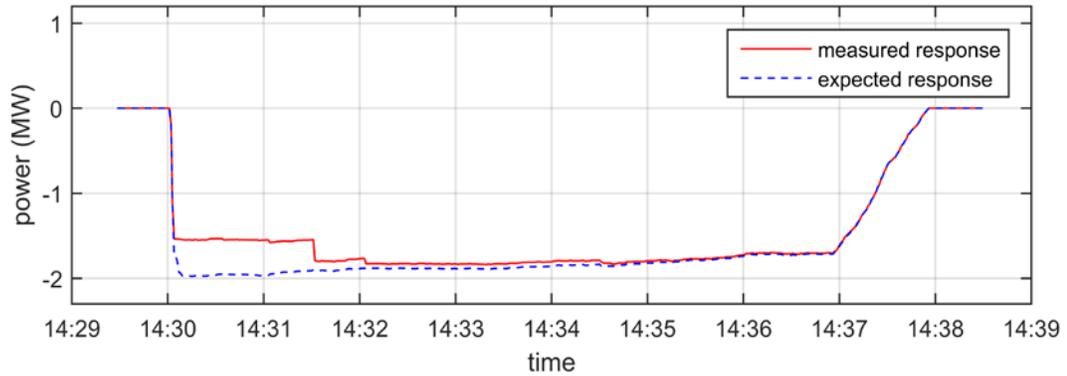


Figure 34 Response in MW for Case 2-4 with the under-frequency event.

In Figure 35, we show the voltage profile for the case with both voltage lockout and voltage sorting activated. Again, the snapshot of the voltage is taken one minute after the contingency. It is clear that implementing both voltage lockout and sorting in this case has a desirable effect. From the plot we see that the voltage profile is generally lower than the one in Case 1.

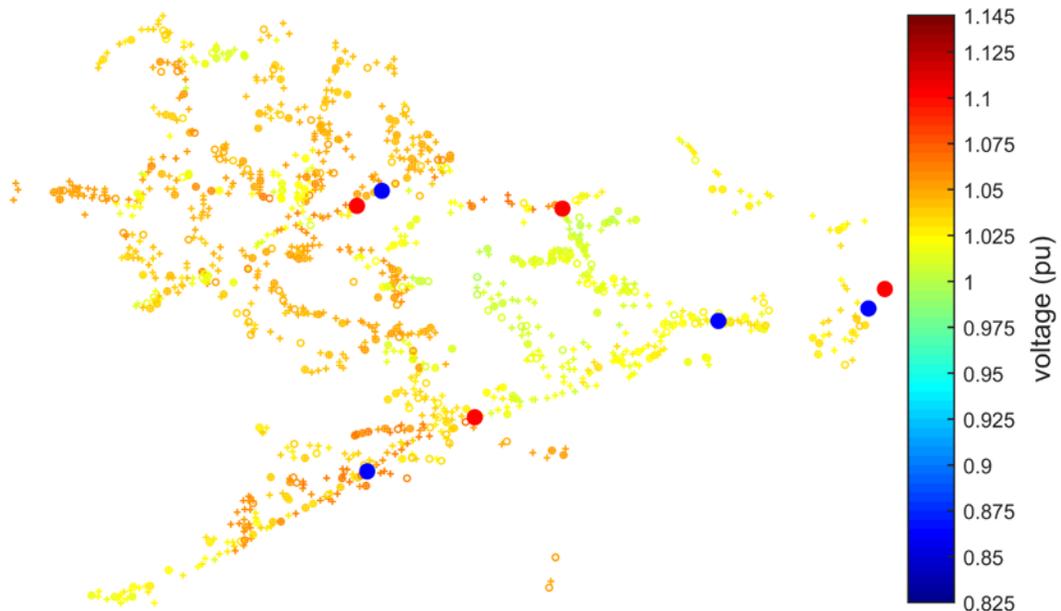


Figure 35 voltage profile for Case 2-4 with under frequency event taken at 14:31. Red dots corresponds to voltage regulators and blue dots are capacitor banks. '+' symbols corresponds to water heater changing state due to the emergency event.

This is also confirmed when looking at the total amount of violations in Table 15, where we see that we do not have any instantaneous voltage violations. We also see a decrease in the number of continuous violations.

Table 15 Violation data for Case 2-4 during under frequency event.

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	29	0	0	0

With the same settings, we will be running Case 2-4 with the over frequency event reported in Case 1. Figure 36, shows the response for Case 2-4 with the over frequency event. In this plot we have both the expected curve based on the supervisor settings and the observed curve. From the figure, it is clear that having devices in voltage lockout is impacting the response from the devices. In this Case we see both a difference in initial response and during the entire event with a difference above 25% from the expected response. From this plot it is also easily seen how the devices stay in lockout for one minute at a time. The reason for the drastic jumps in the response curve lies in the fact that voltage lockout is trying to counteract the undesirable effects of the voltage sorting. Had we increased the voltage lockout period, we might have seen less voltage problems.

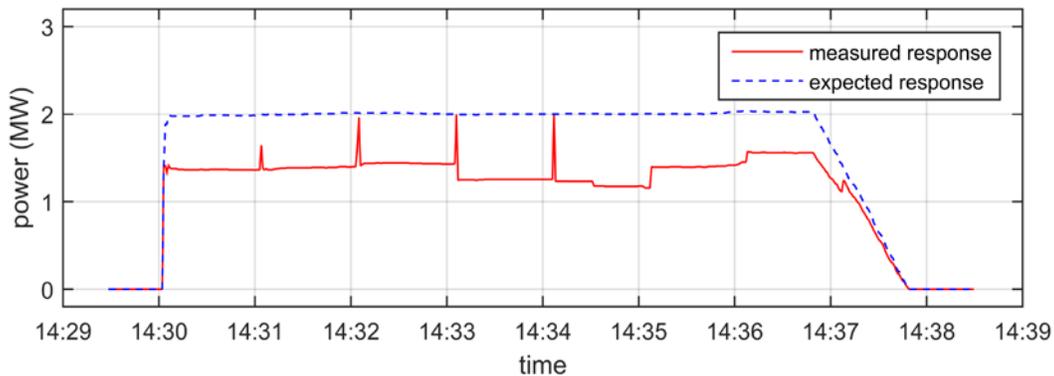


Figure 36 Response in MW for Case 2-4 with over-frequency event.

In Figure 37, the voltage profile for Case 2-4 with the over frequency event is reported. In this, it can be clearly seen that doing voltage lockout and sorting is better than the Case 1 without any of the controls.

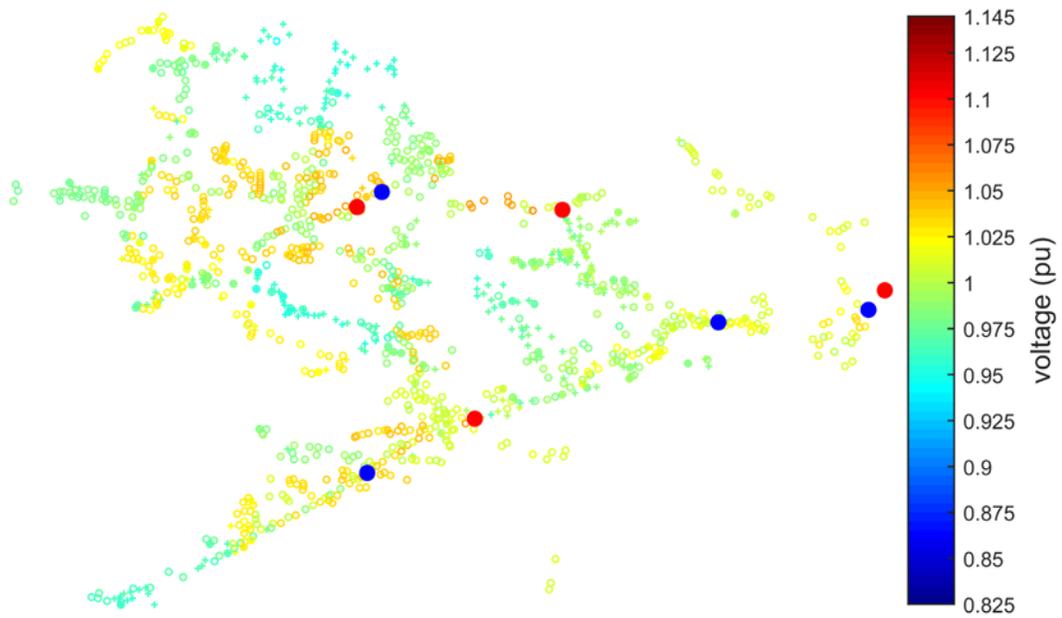


Figure 37 Voltage profile for Case 2-4 with over-frequency event at 14:31. Red dots correspond to voltage regulators and blue dots are capacitor banks. '+' symbols correspond to water heaters changing state due to the emergency event.

However, as seen in Table 16, it does not eliminate all voltage violations. From this table, we still see 78 instantaneous high voltage violations along with 300 continuous high voltage violations.

Table 16 Violation data for Case 2-4 during the over-frequency event.

Voltage in pu	Continuous Voltage Violation (5min)		Instantaneous Voltage Violation (1s)	
	High Voltage (>1.05)	Low Voltage (<0.95)	High Voltage (>1.10)	Low Voltage (<0.90)
Violation count	300	0	78	0

5 Conclusions and Future Work

The impacts and effectiveness of using large-scale deployment of Frequency Responsive Loads to control the primary frequency response after an under-frequency event were studied. The WECC model for the 2015 heavy-load summer was used as the test model in PowerWorld. The study looked at how the FRL penetration level, together with two different ways of assigning control frequencies to the water heater population, can mitigate an under-frequency event. The threshold frequency assignment was done in both decentralized and supervised modes. It was concluded that depending on the level of FRLs penetration, with each strategy to assign the threshold frequencies having its own advantages. While at low penetration levels, the decentralized method is definitely less efficient, increasing the penetration level could make this method at least as efficient as the supervised one, while eliminating the need for communication between devices and the supervisor. However, as the penetration levels of the FRLs increase, the decentralized methodology could potentially lead to overshoot and even instability, depending on the type of contingency. On the contrary, having the curtailment thresholds calculated by the supervisor at fixed time intervals based on the current state of the system (including possible contingency), standard requirement and the available controllable load, the curtailed load is modulated so that only the necessary amount of load is dropped in order to arrest the under-frequency event.

The impacts of deploying frequency responsive load controllers on a distribution system was also demonstrated. In all the studies, a relatively large contingency and very high local penetration of responsive devices was assumed. This is not true for all systems, as many distribution systems would not necessarily support 2 MWs of electric water heater load. The IEEE 8500-Node Test System was used, due to its existing voltage control difficulties, to highlight some of the most extreme impacts that could be seen. Additionally, this particular system only looked at one two-hour window in time with one specific load pattern. The dynamics of the system will change depending on the amount of load available, which is in turn driven by time-of-day, day-of-week, and the season. In extreme cases, the control system may look a lot a “cold load pickup” problem, which tends to synchronize loads after an outage and can cause significant operational issues [26]. Some potential directions of future work are discussed next.

In the results reported, some examples of undesirable behavior under extreme penetration levels were presented. However, if the supervisory droop setting is greater, meaning the devices are less responsive to frequency deviations, the impacts on the distribution system are reduced. How the distribution system is constructed and operated plays a much larger role. Circuits with fewer existing voltage control devices may potentially have fewer deviations. Further investigations into different distribution system controls, settings and topologies are needed to fully quantify the distribution system impacts. In addition, the length of time of impact (much less than two hours) and infrequency of events may mean that distribution utilities can live with this level of impact or require minimal upgrades and/or compensation to support certain levels of penetration. Individual utility practices, circuit topologies, and operational controls will impact this decision making.

Two algorithms were designed to mitigate the voltage impacts on the distribution system. The first one implemented was a local algorithm at the device level monitoring the voltage deviation from

nominal. This is a very simple and crude algorithm, and the algorithm could have an impact on other existing controls on the distribution system, such as voltage regulator and capacitor bank control. It also reduced the effectiveness of the transmission-level frequency control system, discounting the benefits of the system and potentially promising resources that are not available. Further development and testing is needed to analyze the implications of implementing this voltage control algorithm.

Secondly, we implemented two sorting algorithm at the supervisory control level. Both algorithms proved in some cases to cause an undesirable effect of topologically grouping the responsive devices, but ensured that the transmission-level control goals were met. It also created an issue of fairness among devices; voltage sorting could potentially ask more of devices located in a specific region to respond more often. These algorithms treated the transmission and distribution control objectives as independent solutions; future work may investigate a more integrated transmission and distribution control system.

In this work, we always assumed that the supervisor was calculating new frequency thresholds right before the contingency occurred. The droop-like curve is most accurate at the beginning of the threshold re-calculation period. In our results, we get the highly aligned responses to the expected response according to the droop setting of the supervisor. In future work the impacts of the control period for the supervisor should be investigated.

6 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] D. J. Hammerstrom, J. Brous, D. P. Chassin, G. R. Horst, R. Kajfasz, P. Michie, T. Oliver, T. A. Carlon, C. Eustis, O. M. J"arvegren, W. Marek, R. L. Munson, and R. G. Pratt, "Pacific Northwest GridWise testbed demonstration projects Part II. Grid Friendly Appliance project," Pacific Northwest National Laboratory, Technical Report, PNNL-17079, Oct. 2007.
- [4] J. Smith, M. Milligan, E. DeMeo, and B. Parsons, "Utility wind integration and operating impact state of the art," *IEEE Transaction on Power Systems*, vol. 22, no. 3, pp. 900–908, Aug. 2007.
- [5] Y. Makarov, C. Loutan, J. Ma, and P. de Mello, "Operational impacts of wind generation on California power systems," *IEEE Transaction on Power Systems*, vol. 24, no. 2, pp. 1039–1050, May 2009.
- [6] J. Lian, Y. Sun, L. D. Marinovici, and K. Kalsi, "Improved controller design of Grid Friendly Appliances for primary frequency response," in *Proceedings of 2015 IEEE PES General Meeting*, Denver, CO, Jul. 2015.
- [7] N. Lu and D. J. Hammerstrom, "Design consideration for frequency responsive Grid Friendly Appliances," in *Proceedings of 2005/2006 IEEE PES Transmission and Distribution Conference and Exhibition*, Dallas, TX, May 2006, pp. 647–652.
- [8] http://smartgrid.epri.com/doc/amidata_presentation.pdf
- [9] 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.
- [10] PowerWorld Simulator software, <http://www.powerworld.com>
- [11] 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>
- [12] GE PSLF software, <http://www.geenergyconsulting.com/practice-area/software-products/pslf>
- [13] "WECC MVWG Load Model Report ver.1.0," Western Electricity Coordinating Council, June 12, 2012
- [14] "WECC Composite Load Model (CMPLDW) Specifications," Western Electricity Coordinating Council, December 4, 2014
- [15] "WECC Composite Load Model with DG Specifications," Western Electricity Coordinating Council, February 27, 2015
- [16] "Western Electricity Coordinating Council Off-Nominal Frequency Load Shedding Plan", May 2011
- [17] Standard BAL-003-1 — Frequency Response and Frequency Bias Setting, Version 1, February 7, 2013
- [18] IEEE PES Distribution System Analysis Subcommittee's Distribution Test Feeder Working Group, Online at: <http://ewh.ieee.org/soc/pes/dsacom/testfeeders/>
- [19] R. F. Arritt, R. C. Dugan, "The IEEE 8500-Node Test Feeder", *Proceedings of the 2010 IEEE PES Transmission and Distribution Conference and Exposition* New Orleans, 2010.

- [20] K. Schneider and J. Fuller, "Voltage control devices on the IEEE 8500 node test feeder", Proceedings of the *2010 IEEE PES Transmission and Distribution Conference Expo.*, pp.1 -6, 2010.
- [21] K. Schneider, J. Fuller and D.P. Chassin, "Multi-State Load Models for Distribution System Analysis", *IEEE Transactions on Power Systems*, vol.26, no.4, pp.2425-2433, Nov. 2011.
- [22] Z.T. Taylor, K. Gowri and S. Katipamula, "GridLAB-D Technical Support Document: Residential End-Use Module Version 1.0", Pacific Northwest National Laboratory, Technical Report, PNNL-17694, Richland, WA, July 2008.
- [23] J.C. Fuller, N. Prakash Kumar and C.A. Bonebrake, "Evaluation of Representative Smart Grid Investment Grant Project Technologies: Demand Response", Pacific Northwest National Laboratory, Technical Report for the US Department of Energy, 2012.
- [24] W. Marion, K. Urban, User's Manual for TMY2s, National Renewable Energy Laboratory, June 1995, Available: http://rredc.nrel.gov/solar/old_data/nsrdb/tmy2/.
- [25] PacifiCorp Engineering Handbook, Vol. 1, Part C, Section 5.1, Online at: https://www.pacificpower.net/content/dam/pacific_power/doc/Contractors_Suppliers/Power_Quality_Standards/1C_5_1.pdf.
- [26] K. Schneider, E. Sortomme, M. T. Miller, S. S. Venkata, and L. Ponder, "Evaluating the Magnitude and Duration of Cold Load Pick-up using Multi-State Load Models," accepted *IEEE Transactions on Smart Grid*.

Distribution

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

1 Wei Zhang
404 Drees Labs,
2015 Neil Ave.,
Columbus, OH, 43210
(PDF copy)

7 Local Distribution

Pacific Northwest National Laboratory

Karanjit Kalsi (PDF)

Jianming Lian (PDF)

Laurentiu Marinovici (PDF)

Marcelo Elizondo (PDF)

Jacob Hansen (PDF)

Jeff Dagle (PDF)

Robert Pratt (PDF)

Jason Fuller (PDF)

Tess Williams (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