

Bulk Electric Power System Risks From Coordinated Edge Devices

RICHARD KENYON^{1,2} (Student Member, IEEE), JEFF MAGUIRE¹, ELAINA PRESENT¹,
DANE CHRISTENSEN¹ (Member, IEEE), AND BRI-MATHIAS HODGE^{1,2} (Senior Member, IEEE)

¹National Renewable Energy Laboratory (NREL), Golden, CO 80401 USA

²Department of Electrical, Computer and Energy Engineering, University of Colorado, Boulder, CO 80309 USA

Corresponding author: Dane Christensen (dane.christensen@nrel.gov)

This work was supported by the Laboratory Directed Research and Development at NREL, authored by NREL operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract DE-AC36-08GQ28308.

ABSTRACT As smart load adoption grows on the electric power system, potential for losing load diversity increases, possibly in ways that impact system stability. Cloud computing resources are able to coordinate large amounts of behind-the-meter loads and resources. Inadvertent or malicious actions could potentially result in gigawatts of load, distributed across large regions, acting nearly simultaneously. We study the resulting impacts of such a perturbation, which were previously recognized, with improved fidelity and granularity using a physically-based power system and demand model. The ResStock tool was used to calculate residential air conditioning load at more than 3,000 locations across the Western Interconnection, corresponding in time to heavy summer and light spring loading. Under an assumption that one cloud platform managed smart thermostats controlling 10%, 15%, or 20% of residential air-conditioning, calculated load steps could be injected into Positive Sequence Load Flow dynamic simulations. These load-driven effects were coupled with two classes of distributed generation ride-through to evaluate the potential for cascading outages. We found frequency deviations in the spring case far exceed the credible contingency event, leading to widespread distributed generation loss, while voltage depressions during the summer loading lead to widespread distributed generation loss and system separation.

INDEX TERMS Bulk power system, residential air conditioning, coordinated command and control, cloud connected, distributed generation.

I. INTRODUCTION

THE RAPID adoption of connected end-use loads presents an opportunity to reach previously unachievable levels of energy efficiency and dynamic demand integration with grid operations. As markets adapt to Federal Energy Regulatory Commission rulings 745 [1], 841 [2], and 2222 [3], grid-interactive loads will become more valuable, further accelerating their adoption.

Accompanying the benefits, such devices could have potential vulnerabilities which are exploited to impact electric power system operations. When central cloud computing resources provide command and control access to large numbers of energy-consuming devices, there are numerous new risks to consider. Inadvertent algorithmic glitches or platform errors could cause simultaneous actions across all coordinated devices [4]. Human error, either unintentional or intentional [5] could result in undesired or unexpected responses.

As we rely increasingly on central control platforms and algorithms, legal responsibility for such malfunctions is not well-established [6].

The cloud interface also presents an attack surface for a malicious actor. The compromise of individual edge devices has been widely studied and published, including the penetration of individual devices within the internet of things [7]. An edge device botnet have been used to launch a large-scale denial of service attack on the bulk internet [8]; compromised edge devices were leveraged to pivot into the Ukrainian power grid and force a significant outage [9]; and we are aware that U.S. utilities are increasingly concerned about operational disruptions occurring through their digital networks and are bolstering their cybersecurity postures [10].

Regardless of how it occurs, there is potential risk that simultaneous digital operation of large quantities of grid-edge devices could result in negative impacts on the physical

world. This paper presents scenarios over which today's utilities and grid operators have no control—disruption of the power system by modifying large amounts of aggregate demand.

This work presents analysis of coordinated demand assets—residential air conditioners—to expose a currently unmitigated risk to the U.S. power grid, which may be further exacerbated by distributed generation ride through adherence. Our work extends the state-of-the-art [11] by using three best-in-class models: the Western Interconnection, the U.S. residential building stock, and residential air conditioners.

The contributions of this work are summarized as follows:

- The realistic potential of the coordinated control of cloud connected loads to yield a significant, instantaneous change of system net load is studied.
- The spatial and granular dispersion of residential air conditioning capacity across the western US is established, as is the dispatch at two specific times.
- Dynamic simulations are performed to study a realistic emergent system disturbance: an instantaneous residential demand step comprised of individually small while collectively significant loads.
- The cascading impacts of a collective residential air conditioning load step upon solar photovoltaic distributed generation is investigated depending on the adherent ride through criteria, and a follow on vulnerability is highlighted.

LIST OF ACRONYMS

CMPLDWG	: composite load with generation
DG	: distributed generation
GW	: gigawatt
HS	: heavy summer
LFmD	: PSLF motor loading factor
LSP	: light spring
MVA	: mega-volt-ampere; apparent power
MW	: megawatt
PG&E	: Pacific Gas & Electric
PSLF	: Positive Sequence Load Flow
PST	: Pacific Standard Time
p.u.	: per unit
PV	: solar photovoltaics
PVD1	: PSLF generic solar inverter model
RAC	: residential air conditioner
VS	: volt-sec metric
VSDG	: volt-sec-distributed generation metric
WECC	: western interconnection
WWSIS	: western wind and solar integration study

II. BACKGROUND

A. CLOUD COORDINATION OF RESIDENTIAL AIR CONDITIONER LOAD

A recent report [11] found a significant risk from coordinated residential air conditioners (RACs) when applied to a test distribution feeder and an approximate model of Poland's

grid. Other researchers identified a method of attack on an urban power system by manipulating building occupants with false demand response messages [12]. Some recent work has investigated how bulk grid cybersecurity requirements might be applied to aggregated end-use loads that provide market services [13].

In this paper, we look specifically at loads that may be small individually, but that are operated by one or more cloud computing platforms, including those managed by a third-party aggregator. The example used is cloud-connected RACs in the western United States; as the internet of things grows, other loads—e.g. water heaters, electric vehicles, batteries, connected solar inverters—are likely to comprise similar risk to power grids. There are more than 17 million connected thermostats deployed in North America today (and this is expected to grow to at least 60 million by 2027 [14]). Each of these operates a major load, the RAC, which is the category that dominates utility peak load planning. Of the 118 million households in the United States [15], 90% have air conditioning [16], which means roughly 15% of households (and presumably a significantly higher fraction of RAC demand, roughly 25%) is currently served by connected thermostats.

B. ResStock MODEL

To accurately determine the energy consumption of RACs and their maximum potential energy consumption, we used ResStock [17]—a tool for simulating the U.S. single-family housing stock. It provides sampled probability distributions for each feature of U.S. homes along with scaling multipliers, which together form a representative simulation model of a region's residential buildings. A large number of samples (typically 350,000) are used to ensure that the results are representative of the considerable variety in the housing stock. Annual simulation of the housing stock using EnergyPlus [18] was performed using actual 2012 weather data for 216 sites across the contiguous United States. This diversified approach allows for a geographically-accurate representation of, among other things, the RAC energy consumption and the distribution of RAC types and sizes. ResStock outputs include the RAC energy consumption each hour.

C. WESTERN INTERCONNECTION MODEL

The U.S. western interconnection (herein annotated WECC) is the synchronous power system serving 14 states in the western United States, as partitioned by the dashed line in Figure 1. The WECC summer peak demand as of 2019 is approximately 150 gigawatts (GW), which consists of around 80 million electricity customers connected through nearly 120,000 miles of transmission lines [19].

The Western Wind and Solar Integration Study (WWSIS) Phase III [20] investigated the transient stability of the WECC with a large share of renewable energy sources. Ten-year planning cases from 2012 were populated with high shares wind, utility and distributed solar photovoltaics (PV), and concentrating solar power. A light spring (LSP) case—April 4, 2012, at 7:50 a.m. Pacific Standard Time (PST)—and

TABLE 1. Pertinent parameters of WWSIS Light Spring (LSP) and Heavy Summer (HS) cases.

Parameter		LSP (GW)	HS (GW)
Load	Total Load	117.4	192.1
	Modeled with CMPLDWG	95.1	143.9
Generation	Wind	27.0	14.3
	Concentrating Solar	8.5	6.6
	Distributed PV	7.3	9.4
	Utility-Scale PV	10.2	11.2
	Other	65.3	154.5

a heavy spring (HS)—August 25, 2012 at 1:20 p.m. PST—from WWSIS were used in this study. The dynamic model used for the WWSIS study, which was built within the Positive Sequence Load Flow (PSLF) simulation environment [21] on top of the current WECC dynamic model of the time, was used for this work. The authors have extensive experience with this PSLF WECC model [22]–[24]. Table 1 provides pertinent generation and load data—including the composite load with generation (CMPLDWG) dynamic model representation—for these two cases.

III. METHODOLOGY

A. CONTROLLABLE RESIDENTIAL AIR CONDITIONING LOAD ESTIMATION

The initial and maximum controllable RAC loads were calculated as follows. ResStock was used to determine operating (online) RAC load accounting for diversity. The max RAC loads are calculated using the RAC-rated equipment cooling capacities from the ResStock sampling routines, the known outdoor conditions for that hour from the actual meteorological year weather file, and the energy input ratio and capacity coefficients [25], which are the same coefficients used in the ResStock model.

To increase the geographic granularity, RAC load was grouped by weather file location, building vintage bin, building size bin, heating fuel, and cooling type, and then disaggregated into census tracts by merging these grouped results using information from [26]. The results were then downsampled to include only the WECC territory census tracts.

Census tract loads were assigned to the nearest bus based on latitude and longitude, splitting the load equally among collocated buses. The majority of the loads were allocated to existing RAC motor load in the HS and LSP dynamic models used. However, when changes had to be made to the dynamic model loading to fully represent allocated RAC load, the bus’s real-time RAC load capacity was capped at 40% of the total load at the bus, which maintains an existing peak RAC penetration in the initial dynamic model. Once that capacity was reached, loads were assigned to the next-closest bus that still had available capacity. For the HS case, RAC motors were initialized at some CMPLDWG to accommodate ResStock loading.

This load-to-bus assignment was used to assign the online and max RAC loads for the HS and LSP cases, with a summary of a few of the state totals and systemwide totals

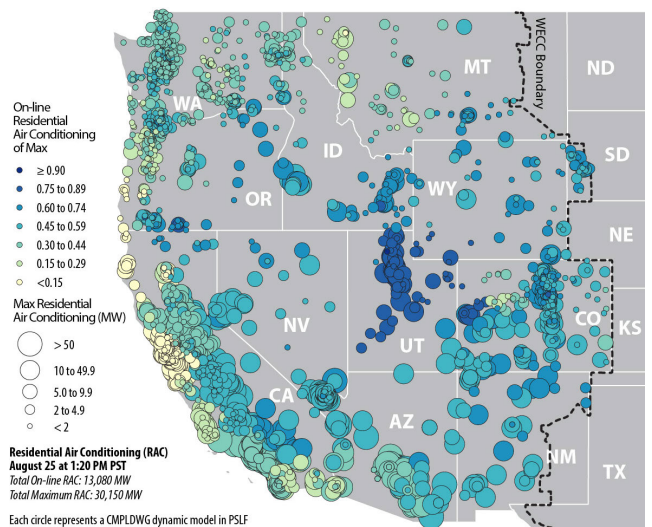


FIGURE 1. ResStock results showing the geographic distribution of online residential air conditioning with respect to the maximum at each CMPLDWG node during the HS case. Dashed line represents WECC boundary.

TABLE 2. Examples of Light Spring (LSP) and Heavy Summer (HS) residential air conditioning real power by region, and WECC total.

Region	LSP RAC (MW)		HS RAC (MW)	
	Online	Max	Online	Max
Arizona	0	5,690	3,410	8,200
California	0	11,340	5,370	14,050
Colorado	0	900	1,010	1,670
Nevada	0	1,390	1,050	1,900
New Mexico	0	680	270	690
WECC Total	0	23,260	13,080	30,150

provided in Table 2. Figure 1 shows the allocated RAC by bus across the WECC system for the heavy summer case, as well as the fraction of RAC online at the load flow time (e.g., 1:20 p.m. PST). During the LSP case, only 10 megawatts (MW) of RAC are online. The associated WECC system image of LSP RAC online is not presented because of a lack of RAC online. The LSP maximums are similar, but less because of thermodynamics, than those for the HS case shown in Figure 1. These data are summarized and broken up by a few of the dominant RAC load states in Table 2.

B. CONTROLLABLE RESIDENTIAL AIR-CONDITIONING IMPLEMENTATION

More than 75% of the load in the HS and LSP PSLF dynamic models is modeled with the CMPLDWG dynamic model; the components of this model are shown in Figure 2. The three motors are defined via generic model parameters to capture distinct, aggregate types of motor load. The static load is voltage- and frequency-dependent, whereas the electronic load is a constant real/reactive power load. Photovoltaic distributed generation (DG) is also included, which reacts according to the PSLF generic solar inverter model (PVD1) model defined in [27]; the voltage and frequency ride-through

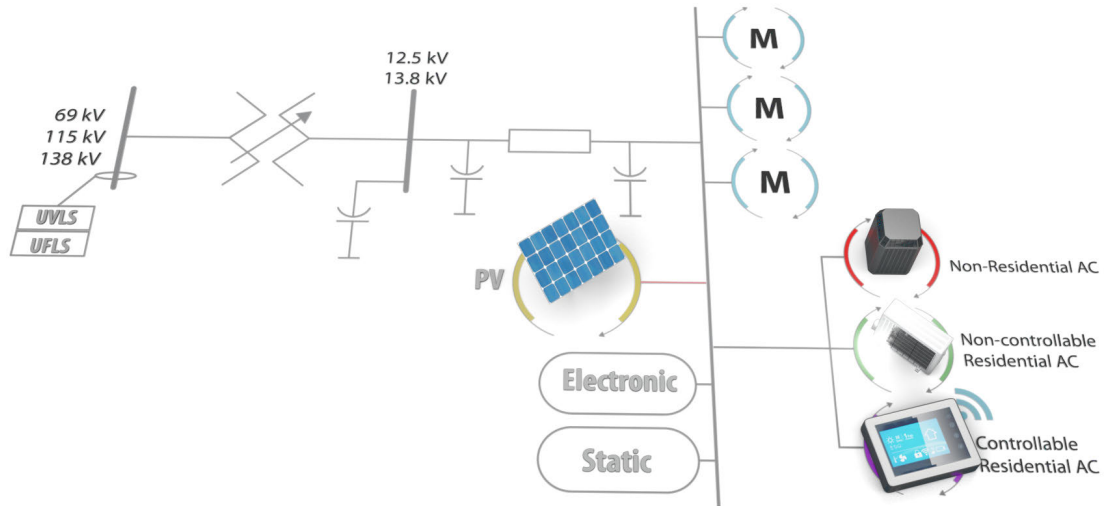


FIGURE 2. The composite load with generation dynamic model, or CMPLDWG, used for more than 75% of all load in the WECC dynamic model. The fourth motor load, broken apart as three types of air conditioning, is a single-phase proxy employed to provide air-conditioning load dynamics to simulations. The controllable load step events are accomplished by adjusting the load fraction on the air-conditioner motor.

characteristics are shown in Figure 6. The DG operates at a unity power factor, with constant power set point. A fourth motor is in the CMPLDWG with the purpose to model the characteristics of single-phase RAC motor load. As shown in Figure 2, for this study, the air conditioning load is partitioned into three types, nonresidential air conditioning, noncontrollable RAC, and controllable RAC. All are still maintained as a single motor in the dynamic model.

To implement the steps in the controllable RAC, PSLF motor loading factor of motor D (LFmD) is adjusted to reflect the commanded step in the controllable AC. This results in a step in the real power consumption of the AC motor D, proportional to the increase from the initialized LFmD (i.e., if LFmD is initialized at 1.0, then changing this to 1.2 during a dynamic simulation will increase the real power consumption by 20%). The loading factor is intended to capture the operating efficiency of the motor at the initialized loading, with changes impacting other values besides the real power consumption. Centering these changes at a unity loading factor minimizes these secondary impacts. This method does not model the in-rush currents that would be exhibited by an air conditioner upon start up [28]. For this reason, these simulations may be considered conservative because the in-rush currents amount to reactive power, leading to larger voltage dips than those caused by a step in real power alone.

Figure 3 shows real power changes for a step off followed by a step to max at the San Mateo CMPLDWG model for three different controllable fractions of the RAC. The implemented fractions in this study are 10%, 15%, and 20%. In Figure 3, we see a consistent quantity of nonresidential AC and noncontrollable RAC that remains constant depending on the controllable fraction. The controllable RAC shows a step to zero real power, followed by a step to the max of the controllable fraction. It is evident that even for only 20% controllable, the real power consumption of the controllable




FIGURE 3. Non-RAC, noncontrollable RAC, and controllable RAC with controllable RAC step off to max, for the heavy summer case at the San Mateo bus.

RACs after stepping to max is more than one-third of the total RAC load at the San Mateo bus.

Figure 4 shows a similar string of controllable RAC steps at the San Mateo bus for the LSP case, but because there is no online RAC at the time of the load flow (7:50 a.m.), there is no RAC to turn off. This results in a rough doubling of the total air conditioner load if only 10% controllable of the RAC units are stepped to max.

C. SIMULATION METHODOLOGY

The aggregate controllable RAC load for each selected fraction of controllable RACs—10%, 15%, and 20%—is shown in Table 3. Note that the total load for the LSP case is 117 GW,




FIGURE 4. Non-RAC, noncontrollable RAC, and controllable RAC with controllable RAC step off step to max, for the light spring case at the San Mateo bus.

TABLE 3. Systemwide Controllable Light Spring (LSP) and Heavy Summer (HS) Residential Air-Conditioning (RAC) totals.

Controllable (%)	LSP RAC (MW)		HS RAC (MW)	
	Online	Max	Online	Max
10	0	2,330	1,310	3,020
15	0	3,490	1,960	4,520
20	0	4,650	2,620	6,030

whereas the total load for the HS case is 192 GW. The result is that the frequency nadir is expected to be lower for the LSP case for the same size load change as compared to the HS case. For a baseline comparison, the system frequency responses for varying controllable fractions are compared against the frequency response following the WECC credible contingency of a double Palo Verde unit trip equating to a loss of 2,750 MW of generation.

Load shedding driven by either under- or over-frequency power system controllers is not modeled in these dynamic simulations in order to preserve a relative impact, between the controllable fractions. Traces representing the frequency load shed values are included for comparison in the time series plots when the system frequency approaches or crosses these limits. These traces are constructed using the Standard PRC-006 criteria as outlined in [29] and [30].

For all simulations an averaged system frequency is presented as a time series. The system frequency is calculated as a rated mega-volt-ampere (MVA) weighted average of all synchronous generator frequencies across the system. Additionally, the volt-sec (VS) and volt-sec-distributed generation (VSDG) metrics as introduced in [31], and exercised in [22]–[24], are calculated for each of the simulations. The VS metric integrates the deviation from a set voltage range of all buses in the system. The VSDG metric is similar, except it weights the deviation by the quantity of DG on the

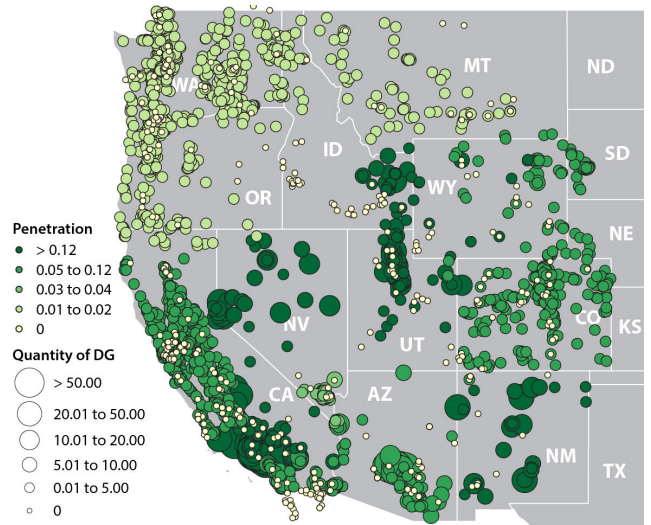


FIGURE 5. Distributed PV located throughout the WECC system, capacity in MW, and penetration as a proportion of native load. Each circle represents the generation simulated with the CMPLDWG dynamic model, at its nodal location.

respective bus. This further highlights voltage fluctuations at DG-containing buses. The running sums and final scores provide valuable voltage impact assessment metrics.

For this study, the voltage deviation boundaries applied to the VS/VSDG metrics are for voltages below 0.95 per unit (p.u.), and above 1.05 p.u. Each dynamic simulation is started and run in steady state for 1 second before the one of the following control schemes are implemented:

- HS: Controllable RAC step off at 1 second
- HS: Controllable RAC step off at 1 second, step to max at 40 seconds
- HS and LSP: Controllable RAC step to max at 1 second
- HS and LSP: Controllable RAC step to max at 1 second, step off at 40 seconds

Note that the capability of a RAC unit to turn on and off, or vice versa, in rapid succession, was not investigated explicitly for this work.

IV. CASCADING IMPACTS ON DISTRIBUTED GENERATION

More than 9 GW of DG is online throughout the system during the HS case, and more than 7 GW for the LSP case. Figure 5 shows the location of DG for the HS case throughout the WECC, with each circle representing a CMPLDWG in the dynamic model, the circle size its capacity in megawatts, and shading the penetration as a proportion of the simultaneous native load. Although these dispatches were based on projections, California alone has an estimated 8 GW of behind-the-meter DG in the form of PV as of 2019 [32]. Thus, the magnitude of DG located in the power system models used is a reasonable representation of the contemporary generation portfolio, although the locational aspect is less easily validated.

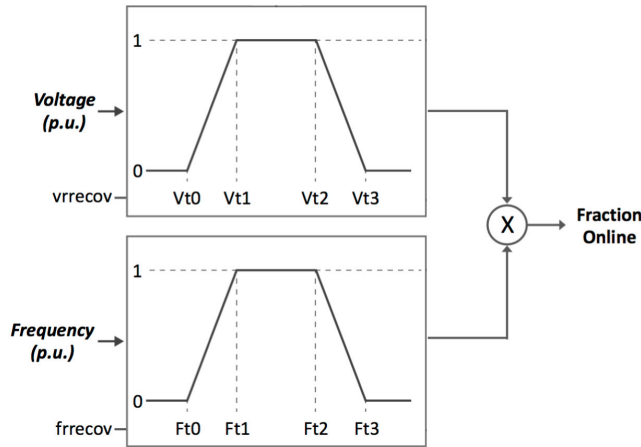


FIGURE 6. Photovoltaic distributed generation voltage and frequency ride through management with the PVD1 dynamic model.

TABLE 4. Distributed generation ride-through parameters.

Parameter	WWSIS	PG&E
vrrecov	1.0	1.0
Vt0	0.7	0.88
Vt1	0.8	0.9
Vt2	1.1	1.1
Vt3	1.2	1.2
frrecov	1.0	0.0
Ft0	50.0	59.5
Ft1	57.0	59.7
Ft2	63.0	60.3
Ft3	65.0	60.5

During frequency or voltage disturbances, the output is scaled according to the implemented ride-through criteria; this process is shown graphically in Figure 6. The defining characteristics are V_{t0} , V_{t1} , V_{t2} , V_{t3} , and V_{rrecov} for voltage, and F_{t0} , F_{t1} , F_{t2} , F_{t3} , and F_{rrecov} for frequency. The four graph points define linear relationships between the fraction online and measured voltage or frequency. From these measurements, the fraction of online DG is derived from the two curves in Figure 6. The recovery fraction defines what fraction is permitted to recover after the frequency/voltage recover. While this model generates a dynamic response, it heavily obscures the temporal ride-through aspects, such as momentary cessation and recovery time, and the individual responses of DG located on a voltage diverse feeder, with a variety of control schemes originating from different manufacturers, which can heavily impact the collective system response [22], [31].

Two sets of ride-through parameters are used for these simulations, which are shown in Table 4 and correspond to Figure 6. The first set was taken from the WWSIS study, where the parameters were chosen to ignore any potential DG loss due to abnormal frequencies or voltages. Recent work by Pacific Gas and Electric (PG&E) [33] used the second set of parameters. For this work, two simulations, one for each parameter set, were performed for each defined RAC controllable fraction. The WWSIS parameters permit an examination of the frequency nadir (lowest frequency following

generation loss) without incorporating the dynamics lost DG, while the PG&E parameters show how the cascading impacts related to DG loss influence the system response.

Note that for the PG&E parameters, the DG lost due to under- or over-frequency is unable to recover. Here, *lost* indicates DG that goes offline during the simulation, and *recovered* indicates DG that returns online before the end of the simulation.

V. RESULTS

The following are the results for the control schemes implemented in the LSP and HS cases. Each control scheme was performed at 10%, 15%, and 20% controllable RAC fractions, as well as each of the ride-through criteria, WWSIS and PG&E, as outlined in Table 4. A total of 38 (6 control schemes \times 3 controllable fractions \times 2 ride-through criteria + 2 Palo Verde contingencies) PSLF simulations were run. A time series of the system frequency is presented as the primary result, with all six relevant traces (or seven with the Palo Verde frequency for the step-to-max schemes) on each plot. The WWSIS ride-through results are shown with solid traces, and the PG&E ride-through with dashed traces. The trace color (green, blue, and magenta) corresponds with the controllable fraction (10%, 15%, and 20%, respectively). Where applicable, the frequency load shed limit is shown as a dotted red trace. Additional result analysis, i.e. VS/VSDG scores, is provided when insightful.

A. LIGHT SPRING

The system frequency traces for the LSP step-to-max simulations are shown in Figure 7. At 10%, no DG is impacted and therefore the two traces for the two different ride-through criteria are identical. The nadirs for the 15%- and 20%-controllable fractions with WWSIS ride-through are below the double Palo Verde unit trip contingency. For the 20% case, the frequency deviates below 59.5 and exhibits oscillations upon recovery. Two frequency relays were triggered, removing a small amount of generation; the frequency oscillations indicate potential system instability. Table 5 provides the DG loss statistics for these three simulations, showing that there was an impact only for the 20% controllable case (a slight loss of 55 MW for the WWSIS ride-through), with a full recovery before the simulation ends.

For the case of 15% and 20% with the PG&E ride-through, all of the DG (7,310 MW) is lost, with none recovering. This loss in generation is nearly three times greater than the double-unit Palo Verde contingency. With a load increase of 3.5 and 4.7 GW in the 15% and 20% cases, respectively, followed by a loss of 7.3 GW of generation, very large frequency deviations are expected. This loss is evident in the frequency traces in Figure 7, as both traces show a precipitous drop in frequency occurring a few seconds after the RAC step to max. Thus, this is a cascaded impact with the RAC step to max influencing the system DG. The 20% case dips into under frequency load shedding territory. Both simulations diverge.




FIGURE 7. Light spring system frequency traces for varying controllable fraction steps to max. 15% and 20% PG&E simulations accompanied by significant (>7 GW) losses of distributed generation.

TABLE 5. Light spring step-to-max distributed generation loss.

Controllable (%)	Distributed Generation Impact (MW)			
	WWSIS Ride-Through		PG&E Ride-Through	
	Lost	Recovered	Lost	Recovered
10	0	0	5	5
15	0	0	7,310	0
20	55	55	7,310	0

Light spring step-to-max-to-off results are similar to the step-to-max case, due to no RAC initially online.

B. HEAVY SUMMER

1) STEP OFF

For the step-off control scheme in the HS case, the controllable fraction of RAC is simply turned off and a frequency rise results, with the peak frequency growing with the magnitude of load reduction. The frequency traces are identical to the first 40 seconds of step off to max simulation in Fig. 10, and are not reproduced here. The frequency deviations are well within any frequency load shedding criteria. No DG is impacted during these simulations, and therefore the PG&E ride-through results perfectly match the WWSIS ride-through results (i.e., 0 MW lost and 0 MW recovered for all six simulations).

2) STEP TO MAX

The HS step-to-max simulations present a few unexpected frequency responses, although none deviate significantly from the nominal. For the WWSIS ride-through, the expected decrease in the nadir results with growing controllable fractions. At 20%, the frequency response is nearly the same as that for the double Palo Verde unit trip contingency, which is reasonable considering that the load increase of 3,400 MW is near the generation loss amount of 2,750 MW. It is likely that the spatial dispersion of the RAC is partly responsible for why the 20% nadir is not lower (even though it exceeds the

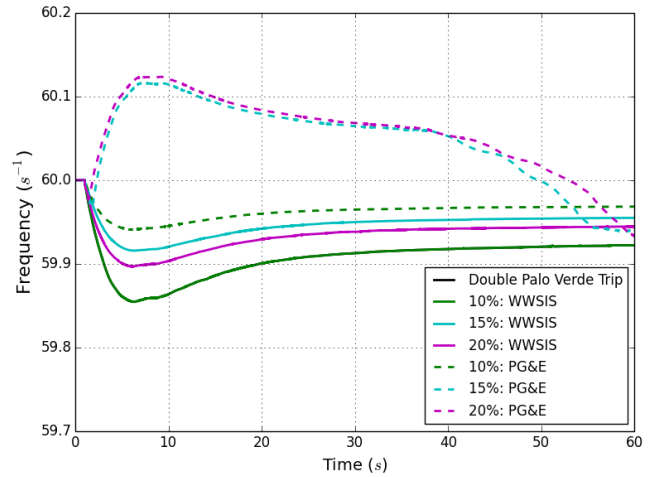


FIGURE 8. Heavy summer system frequency traces for varying controllable fraction steps to max. Note that the double Palo Verde unit trip trace nearly overlays the 20% WWSIS.

TABLE 6. Heavy summer step-to-max distributed generation loss.

Controllable (%)	Distributed Generation Impact (MW)			
	WWSIS Ride-Through		PG&E Ride-Through	
	Lost	Recovered	Lost	Recovered
10	5	5	95	5
15	95	5	1710	1320
20	105	5	1730	325

Palo Verde contingency by 650 MW), as was discovered in the WWSIS study itself when comparing a dispersed DG trip versus a double Palo Verde unit trip at the same power loss [20]. The WWSIS ride-through criteria results in a maximum loss of 105 MW of DG with 95 MW recovering, due to voltage deviations.

The PG&E ride-through yields a frequency nadir that is actually higher for the 10% simulation, and in fact the frequency responses are inverted for the 15% and 20% simulations. Here, the VS metric can help to understand what is occurring. The low-voltage (below 0.95 p.u.) VS and VSDG scores are presented in Table 7. Because these six simulations are the same time length, the same controllable fraction VS and VSDG scores are comparable.

The PG&E score is larger, indicating that either a larger number of buses experienced similar voltage deviations, or a similar number of buses experienced larger voltage deviations. The larger VSDG scores in Table 7 correlate with larger losses in DG as shown in Table 6. The growing scores indicate much larger and/or more time-persistent voltage depressions as compared to the WWSIS simulations.

Figure 9 is a distribution of voltages for buses with voltages below 0.95 p.u., for the PG&E 20% controllable case. The black trace represents the quantity of buses with voltages below 0.95 p.u. captured by the distribution, whereas the red scale traces provide percentiles of the distribution. It is evident that more than 600 buses (out of 21,000 total) have voltages below 0.75 p.u. for more than 40 seconds. We see that as the simulation continues, the voltages improve

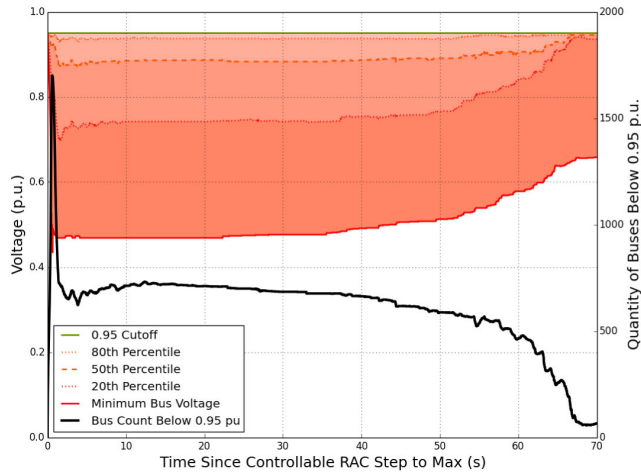


FIGURE 9. Low-voltage distribution and bus count for the controllable RAC step to max. A large voltage suppression is evident, resulting in significant load/motor tripping.

TABLE 7. Heavy summer step-to-max Volt-Sec(-DG) scores.

Controllable (%)	Volt-Sec and Volt-Sec-DG Scores			
	WWSIS Ride-Through VS	WWSIS Ride-Through VSDG	PG&E Ride-Through VS	PG&E Ride-Through VSDG
10	240	940	550	3,670
15	700	4,150	8,810	29,600
20	910	4,710	10,720	33,090

significantly. Further investigation shows that the total static load in the system, which is voltage dependent, decreased significantly following the large loss in DG, which occurred just after the RAC step. Additionally, a large quantity of motor load was temporarily tripped offline because of low voltages. This explains the frequency rise, because although the RAC stepped up considerably, and generation was lost, the total decrease in load (RAC – (static + motor load)) exceeded this DG loss, resulting in the frequency rise. Thus, the RAC step to max causes an initial voltage decrease, causing a loss in DG, which causes a further voltage decrease that eventually results in a net decrease in load because of the voltage-dependent static load and motor tripping. The WWSIS study noted that a loss in real power DG comes with a fourfold increase in reactive power load [20]. Thus, such a significant voltage depression is not surprising following the loss of 1.7 GW of DG.

3) STEP TO OFF TO MAX

If all controllable RAC units are off, and then stepped on, the resultant load increase is roughly double that of turning all controllable RAC units on with half online at the start of the simulation. The system frequencies are shown in Figure 10. The first 40 seconds are identical to the step-off results, whereas the frequency decline following the step to max is larger than just the step-to-max simulations. For the PG&E simulations, the DG losses are similar between the step-to-max, and step-off-to-max simulations for the 10% and 15% cases, although far less recovers in the 15% case. For the 20% case, all DG is lost at 9,680 MW and the system separates,




FIGURE 10. Heavy summer system frequency for varying controllable fraction steps to off to max. System separates for the 20% controllable PG&E case.

TABLE 8. Heavy summer step-to-off-to-max distributed generation losses.

Controllable (%)	Distributed Generation Impact (MW)			
	WWSIS Ride-Through Lost	WWSIS Ride-Through Recovered	PG&E Ride-Through Lost	PG&E Ride-Through Recovered
10	5	5	130	35
15	105	10	1,730	120
20	140	45	9,680	0

as evident by the sharp (ambiguous) frequency increase in Figure 10.

C. RESULTS SUMMARY

In the light spring conditions, when the total loading on the system is far less than the heavy summer case and the online RAC is nearly zero, large frequency swings occurred following controllable RAC steps. With the PG&E ride-through criteria, all DG was lost following frequency deviations beyond the ride-through criteria in the 15% and 20% controllable cases, which resulted in larger frequency deviations into frequency load shed territory followed by simulation divergence. For the heavy summer conditions, the frequency swings were smaller as compared to the spring conditions for two reasons: roughly half of the RAC was online at the start of the simulations, resulting in a smaller amount of a load swing, and the total load on the system was much higher. Regardless, it was found that large-voltage depressions occurred during steps to max, which caused large DG losses that exacerbated the system response. In a few of the cases, the voltage depressions were significant enough to cause an overall decrease in load because of the decrease in voltage-dependent load and some low-voltage motor tripping. For the step-to-off-to-max 20% controllable PG&E ride-through criteria simulation, all 9 GW of DG were lost and the system separated.

VI. CONCLUSION

Within the assumptions and constraints of the methods used, our results indicate that centrally-coordinated load can pose

a catastrophic risk to the bulk power grid, and multiple bulk power system risks can result if this load is synchronized in a step-wise manner. In particular, it was found that even modest fractions of controllable residential air conditioner load, during times when this load is not in use such as the Spring, could create a load/generation imbalance far exceeding the WECC credible contingency, resulting in system instability. These risks may be compounded depending on the distributed generation ride through response, which might lead to widespread disconnection due to the local voltage and frequency response to the synchronized load perturbations.

VII. FUTURE WORK

There are numerous opportunities to expand and enhance this work. Motor inrush from air conditioners was not considered, but will have an additional deleterious impact under some conditions mentioned. The dynamic model should be expanded to include additional bulk power grid dynamics. Our work should be expanded to include cloud vulnerabilities associated with other grid-edge devices—such as electric vehicles, rooftop solar photovoltaics, and batteries. It is also worth considering potential telecommunication limitations on triggering this many internet-connected devices simultaneously, and how network latency might impact these results—if at all.

Lastly, this work only explored power grid impacts that could result from a centrally-controlling cloud. A more comprehensive risk assessment could define the relative benefits of mitigating vulnerabilities at these type of cloud coordinators, versus other potential threat vectors such as utility supervisory control and data acquisition systems. This type of assessment is left for future work.

ACKNOWLEDGMENT

The authors are grateful to our management—especially J. Torres, J. Green, B. Kroposki, C. Kutscher, A. Karagiozis, J. Barnett, and S. Terry—for their strong support and encouragement. B. Roberts and J. Bauer developed graphics for this report. Funding provided NREL under the Laboratory Directed Research and Development program. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

REFERENCES

- [1] *FERC Order No. 745: Demand Response Compensation in Organized Wholesale Energy Markets*, Federal Energy Regulatory Commission (FERC), Washington, DC, USA, Mar. 2011.
- [2] *FERC Order no. 841: Electric Storage Participation in Markets Operated by Regional Transmission Organizations and Independent System Operators*, Federal Energy Regulatory Commission (FERC), Washington, DC, USA, Feb. 2018.
- [3] *FERC Order No. 2222: Participation of Distributed Energy Resource Aggregations in Markets Operated by Regional Transmission Organizations and Independent System Operators*, Federal Energy Regulatory Commission (FERC), Washington, DC, USA, Sep. 2020.
- [4] *Is the Year 2038 Problem the New Y2K Bug?* The Guardian. Accessed: 2014. [Online]. Available: <https://www.theguardian.com/technology/2014/dec/17/is-the-year-2038-problem-the-new-y2k-bug>
- [5] *What's Behind VW's Downfall and Secretive Emissions Deals*. Scientific American. Accessed: 2015. [Online]. Available: <https://www.scientificamerican.com/article/what-s-behind-vw-s-downfall-and-secretive-emissions-deals/>
- [6] M. C. Elish, "Moral crumple zones: Cautionary tales in human-robot interaction," *Engaging Sci., Technol., Soc.*, vol. 5, p. 40, Mar. 2019.
- [7] G. Hernandez, O. Arias, and D. Buentello, "Smart nest thermostat: A smart spy in your home," in *Proc. Black Hat Conf.*, 2014, pp. 1–8.
- [8] M. Lyu, D. Sherratt, A. Sivanathan, H. H. Gharakheili, A. Radford, and V. Sivaraman, "Quantifying the reflective DDoS attack capability of household IoT devices," in *Proc. 10th ACM Conf. Secur. Privacy Wireless Mobile Netw.*, Jul. 2017, pp. 46–51.
- [9] *Analysis of the Cyber Attack on the Ukrainian Power Grid—Defense Use Case*, E-ISAC, Washington, DC, USA, Mar. 2016. [Online]. Available: https://ics.sans.org/media/E-ISAC_SANS_Ukraine_DUC_5.pdf
- [10] R. J. Campbell, "Electric grid cybersecurity," Congressional Res. Service, Washington, DC, USA, Tech. Rep. R45312, 2018. [Online]. Available: <https://crsreports.congress.gov/product/pdf/R/R45312>
- [11] S. Soltan, P. Mittal, and H. V. Poor, "BlackIoT: IoT botnet of high wattage devices can disrupt the power grid," *Proc. 27th USENIX Secur. Symp.*, Aug. 2018, pp. 15–32.
- [12] G. Raman, J. C.-H. Peng, and T. Rahwan, "Manipulating Residents' behavior to attack the urban power distribution system," *IEEE Trans. Ind. Inform.*, vol. 15, no. 10, pp. 5575–5587, Oct. 2019.
- [13] D. Christensen, M. Martin, E. Gantumur, and B. Mendrick, "Risk assessment at the edge: Applying NERC CIP to aggregated grid-edge resources," *Electr. J.*, vol. 32, no. 2, pp. 50–57, Mar. 2019.
- [14] P. Leuschner and C. Talon, "Market data: Advanced thermostats residential smart and communicating thermostats for single-family and multifamily buildings: Global market analysis and forecasts," Navigant, Boulder, CO, USA, Tech. Rep., Dec. 2018.
- [15] *2018 Annual Energy Outlook*, U.S. Energy Inf. Admin., Washington, DC, USA, 2018.
- [16] *The Future of Cooling: Opportunities for Energy-Efficient Air Conditioning*, Int. Energy Agency, Paris, France, 2018. [Online]. Available: <https://www.iea.org/futureofcooling/>
- [17] E. Wilson, C. Christensen, S. Horowitz, J. Robertson, and J. Maguire, "Energy efficiency potential in the US single-family housing stock," NREL, Golden, CO, USA, Tech. Rep. NREL/TP-5500-68670, 2017. [Online]. Available: <https://www.nrel.gov/docs/fy18osti/68670.pdf>
- [18] D. Crawley, L. Lawrie, C. Pedersen, and F. Winkelmann, "EnergyPlus: Energy simulation program," *ASHRAE J.*, vol. 42, no. 4, pp. 49–56, 2000. [Online]. Available: <https://energyplus.net/>
- [19] *State of the Interconnection*, Western Electr. Coordinating Council, Salt Lake City, UT, USA, 2019.
- [20] N. W. Miller, M. Shao, S. Pajic, R. D'Aquila, and G. E. Management, "Western wind and solar integration study phase 3—frequency response and transient stability," GE Energy, Schenectady, NY, USA, Tech. Rep. NREL/SR-5D00-62902, 2014. [Online]. Available: <https://www.nrel.gov/docs/fy15osti/62906.pdf>
- [21] GE Energy Consulting. *PSLF | Transmission Planning Software | GE Energy Consulting*. Accessed: Jan. 2, 2020. [Online]. Available: <http://www.geenergyconsulting.com/practice-area/software-products/pslf>
- [22] R. W. Kenyon and B. Mather, "Simulating distributed energy resource responses to transmission system-level faults considering IEEE 1547 performance categories on three major WECC transmission paths," NREL, Golden, CO, USA, Tech. Rep. TP-5D00-73071, 2020. [Online]. Available: <https://www.nrel.gov/docs/fy20osti/73071.pdf>
- [23] M. Bossart, R. W. Kenyon, D. Maksimovic, and B.-M. Hodge, "The effect of power electronic loads on western interconnection stability," in *Proc. IEEE Power Energy Soc. Gen. Meeting (PESGM)*, Aug. 2020, pp. 1–5.
- [24] R. W. Kenyon, B. Mather, and B.-M. Hodge, "Coupled transmission and distribution simulations to assess distributed generation response to power system faults," *Electr. Power Syst. Res.*, vol. 189, Dec. 2020, Art. no. 106746.
- [25] D. Cutler, J. Winkler, N. Kruis, C. Christensen, and M. Brandemuehl, "Improved modeling of residential air conditioners and heat pumps for energy calculations," NREL, Golden, CO, USA, Tech. Rep. NREL/TP-5500-56354, 2013. [Online]. Available: <https://www.nrel.gov/docs/fy13osti/56354.pdf>

- [26] C. Harris, E. Wilson, and J. Robertson, "ResStock county/household income assignment CHIA," NREL, Golden, CO, USA, 2019. [Online]. Available: <https://data.nrel.gov/submissions/113>, doi: 10.7799/1546055.
- [27] *WECC PV Power Plant Dynamic Modeling Guide*. Western Electricity Coordinating Council. Accessed: 2014. [Online]. Available: <https://www.wecc.org/Reliability/WECC%20Solar%20Plant%20Dynamic%20Modeling%20Guidelines.pdf>
- [28] P. Pourbeik and B. Agrawal, "A hybrid model for representing air-conditioner compressor motor behavior in power system studies," in *Proc. IEEE Power Energy Soc. Gen. Meeting Convers. Del. Electr. Energy 21st Century*, Jul. 2008, pp. 1–8.
- [29] *Automatic Underfrequency Load Shedding*, Standard PRC-006-2, NERC, Atlanta, GA, USA, Aug. 2017. [Online]. Available: <https://www.nerc.com/pa/stand/reliability%20standards/prc-006-2.pdf>
- [30] P. Mackin, "Underfrequency load shedding review group (UFLSRG) report," Western Electr. Coordinating Council, Salt Lake City, UT, USA, Tech. Rep., Jul. 2018. [Online]. Available: <http://www.egr.unlv.edu/~eebag/18%20-%20UFLSRG%20Report%20-%20202018%20July%20-%20post.pdf>
- [31] R. W. Kenyon and B. Mather, "Quantifying transmission fault voltage influence and its potential impact on distributed energy resources," in *Proc. IEEE Electron. Power Grid (eGrid)*, Nov. 2018, pp. 1–6.
- [32] N. A. E. R. Corporation, "Distributed energy resources: Connection modeling and reliability considerations," NERC, Atlanta, GA, USA, Tech. Rep., Feb. 2017. [Online]. Available: https://www.nerc.com/comm/Other/essntrlblysrvcstskfrDL/Distributed_Energy_Resources_Report.pdf
- [33] Pacific Gas and Electric Company, "EPIC 2.05 inertia response impact for DG impact improvement," Pacific Gas Electr., San Francisco, CA, USA, Tech. Rep., Feb. 2019. [Online]. Available: https://www.pge.com/pge_global/common/pdfs/about-pge/environment/what-we-are-doing/electric-program-investment-charge/PGE-EPIC-Project-2.05.pdf



RICHARD (WALLACE) KENYON (Student Member, IEEE) received the double B.A. degree in mathematics and physics from the University of Vermont and the M.S. degree in electrical engineering from the University of Colorado, where he is currently pursuing the Ph.D. degree in electrical engineering. He focuses on power system dynamics and control schemes for stable operation as distributed energy resources grow.



JEFF MAGUIRE received the B.S. degree in mechanical engineering from the Rensselaer Polytechnic Institute, and the M.S. degree in mechanical engineering from the University of Colorado at Boulder. He joined the NREL's Buildings and Thermal Sciences Center, in 2012. His research focuses on the modeling and performance analysis of residential energy demand, analyzing field data, and studying grid-interactive efficient buildings. He specializes in component modeling and performance analysis.



ELAINA PRESENT received the B.S. degree in mechanical engineering from the Massachusetts Institute of Technology, and the M.S. degrees in architecture/building science and in civil and environmental engineering from the University of California, Berkeley. She is currently a Research Engineer with the Buildings and Thermal Sciences Center, NREL. She studies residential energy consumption patterns, how they may change in the future, and the resulting impacts on the evolving energy system.



DANE CHRISTENSEN (Member, IEEE) received the B.S. degree in mechanical engineering from Rice University, and the Ph.D. degree in mechanical engineering from the University of California, Berkeley. He is currently a Principal Engineer for Energy Security and Resilience at the NREL. He studies methods and metrics for evaluating and improving energy system security and resilience at all scales. He also has a keen interest in grid-interactive efficient buildings, energy blockchain, and emerging sensors, analytics, and controls.



BRI-MATHIAS HODGE (Senior Member, IEEE) received the B.S. degree in chemical engineering from Carnegie Mellon University, the M.S. degree in chemical engineering from Åbo Akademi University, and the Ph.D. degree in chemical engineering from Purdue University. He is currently a Chief Scientist with the NREL's Power Systems Engineering Center, an Associate Professor with the Department of Electrical, Computer and Energy Engineering, University of Colorado Boulder, and a Fellow of the Renewable and Sustainable Energy Institute. His expertise is modeling and simulation power systems and ancillary service markets.

...