Fall 12-15-2017

Co-simulation Framework for Power Distribution System Analysis with Humans in the Loop

Victor H. Ayon
University of New Mexico - Main Campus

Candidate

Department

This thesis is approved, and it is acceptable in quality and form for publication:

Approved by the Thesis Committee:

, Chairperson

Co-simulation Framework for Power Distribution System Analysis with Humans in the Loop

by

Victor H. Ayon

A.S., Central Arizona College, 2010
B.S., Mechanical Engineering, University of New Mexico, 2014

THESIS

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Master of Science
Mechanical Engineering

The University of New Mexico

Albuquerque, New Mexico

December, 2017

Dedication

To my beautiful wife, Saydra, for her unconditional love, support, and encouragement. She motivates me and makes me a better person. She is my rock, my best friend, and my soulmate.

"You never fail until you stop trying."
– Albert Einstein

Acknowledgments

I would like to thank my advisor, Professor Andrea Mammoli, for his support and guidance throughout this journey. I would also like to thank my friend and colleague, Matthew Robinson, who was always there to lend a helping hand with support and encouragement.

Co-simulation Framework for Power Distribution System Analysis with Humans in the Loop

by

Victor H. Ayon

A.S., Central Arizona College, 2010
B.S., Mechanical Engineering, University of New Mexico, 2014
M.S., Mechanical Engineering, University of New Mexico, 2017

Abstract

Recent developments in the area of power generation have led to an increased penetration of storage and distributed energy resources (DER) in power distribution systems. As a result, new and enhanced energy management systems will be necessary as the deployment of DERs, as well as the need to control loads, continues to increase in the coming years. Advanced management systems are especially important to achieve resilient power delivery during emergency situations. During a blackout for example, a section of a distribution feeder could island to operate as a microgrid to ensure critical services such as water, food, and medical care remain online. The development of such management systems will also require the ability to integrate human behavior models with power flow simulators, as technological advances lead to more customer-owned devices with the potential to be used to balance the power flow of a feeder. This thesis describes a co-simulation framework that combines a bottom-up residential load generator, a load aggregator for real-time (RT) residential demand response (DR), a utility-scale battery controller, and

the GridLAB-D distribution system simulator. The behavior of a distribution feeder model is analyzed under different scenarios. The model is based on an existing feeder located in Los Alamos, NM, which serves residential customers and a set of critical loads including a hospital, a supermarket, and a water treatment plant. This feeder also hosts a utility-scale solar array and battery storage that are used to operate the feeder as a microgrid. Additionally, a real-time simulation is described in which real-time residential demand response is implemented on a virtual community using the load generator and aggregator. The simulation is part of a project invested in the development of a modern microgrid control system employing a virtual power plant approach and a model predictive controller to optimize the use of resources within a distribution feeder. The capabilities to study power distribution systems with humans in the loop using these platforms is showcased here. Furthermore, their potential as instrumental tools in the development and design of new technology essential to improve grid resilience is also discussed.

Contents

List of Figures

List of Tables

Glossary

DER	Distributed energy resources.
RT	Real-time.
DR	Demand response.
PV	Photovoltaics.
kW	kilo-Watt.
MW	Mega-Watt.
kWh	kilo-Watt-hour.
FNCS	The Framework for Networked Co-Simulation software.
PNNL	Pacific Northwest National Laboratory.
MRI	Mitsubishi Research Institute.
UNM	University of New Mexico.
NEDO	Japan's New Energy and Industrial Technology Development Organization.
HIL	Hardware-in-the-Loop.

Glossary

EV	Electric Vehicle.
VPP	Virtual Power Plant.
MPC	Model Predictive Controller.
BESS	Battery energy storage system.
LADPU	Los Alamos Department of Public Utilities.
GIS	Geographic Information System.
PDF	Probability density function.
HVAC	Heating, ventilation, and air conditioning.
M_S	Thermal mass of the conditioned space by a thermostat.
$T(t)$	Temperature of the conditioned space by a thermostat.
T_a	Ambient temperature.
t	Time.
Q_L	Thermal losses from structure.
Q_R	Thermal gain from air conditioner or heat pump.
K_1	Function of the building's thermal insulation.
COP	Coefficient of performance.
P_{AC}	Power of compressor.
Λ	State function that indicates whether an air conditioner or heat pump is on or off.
T_L	Lower deadband for temperature control.

T_U	Upper deadband for temperature control.
T_{amb}	Outside temperature.
QoS	Quality of service.
TCL	Thermostatically controllable loads.
AC	Air conditioner.
POA	Plane of array.
scp	Secure Copy.
SOC	State of charge.

Chapter 1

Introduction

1.1 Overview

The modernization of the power grid that is currently taking place all over the world has triggered a shift from a century-old paradigm that focuses in centralized power generation and long distance transmission, to a modernized approach that concentrates in localized and distributed power generation. The power grid is composed of three major parts; generation, transmission, and distribution & consumption. Traditionally, power is generated exclusively at power plants and sent to consumers through transmission and distribution lines. However, the power grid is currently undergoing a technological revolution in which power is generated at distributed locations and near to the consumer. Microgrid systems, capable of operating without any support from the main grid have emerged as a result of an increased penetration of distributed energy resources (DER) in distribution level systems. These include utility-scale photovoltaics (PV) and battery energy storage systems capable of delivering 500kW or more, as well as customer-owned rooftop PV, electric vehicles, and other small-scale storage systems including demand response (DR). In the case of

a major disruptive event, such as a hurricane disabling large sections of the power grid, the ability to divide a distribution network into smaller microgrids suggests that local resources could be used to support critical services such as food, water, and medical care, while damaged infrastructure is repaired. For example, the devastating conditions experienced in Puerto Rico after hurricane Maria destroyed most of the island's transmission infrastructure in late September, 2017, could have been significantly reduced by the use of distribution-level microgrids. Water distribution infrastructure, hospitals, and supermarkets could have been given priority to ensure residents had the means to remain hydrated and nourished with the required medical care. The ability to provide critical services during disruptive events indicates that the resilience of the power grid can be improved by considerably reducing the magnitude and duration of these events with proper planning and microgrid resources. Furthermore, DER devices used in microgrids continue to become more economically attractive due to advances in technology in addition to incentives, rebates, and tax credit programs provided by governments. This decrease in cost suggests that the deployment of DERs will continue to increase in the future, and with it, the need for more advanced management systems capable to manage operations from a few kW to several MW.

A well-coordinated deployment of DERs coupled with the development and implementation of control systems could lead to multiple advantages. For instance, inverters could be used to provide local voltage support by adjusting their power factor in order to reduce stress on devices such as load tap changers, potentially increasing their lifetime. Peak loads on substation transformers could be reduced by coordinating the action of storage devices that would lead to deferral of upgrades and replacements. Also, the ability of a distribution feeder to rapidly respond to load adjustment requests could provide frequency support to the transmission grid, thereby reducing the likelihood of blackouts and brownouts. Distribution feeders with the ability to function as microgrids could island sections of a feeder to provide service

to critical loads in the case that a blackout does occur. Also, the load on a distribution feeder could be managed by implementing more sophisticated demand response programs, in which financial incentives are offered to customers willing to reducing or shift their electricity usage during events of high power consumption or other grid constraints.

A well-designed control framework could enable new business models that include and engage customers in the management of residential loads. However, despite current standards and guidelines for the implementation of such framework, it is widely acknowledged that a full deployment is not yet feasible. Further research is needed in the areas of communication, controls, optimization, architecture, and human behavior modeling. While advances have been made in the management and optimization of DERs in the recent years, human behavior research is limited due to the difficulty to test new techniques in realistic situations using real customers and their homes. Especially, studies that examine models of the electric infrastructure in combination with models of how human use electricity are close to non-existent.

This thesis describes the work behind two publications by the author in which a tool to conduct research in the areas of DER management and human behavior models is introduced. The simulation framework provides the ability to model electric distribution systems while considering the effects of human power consumption described by their interactions with household appliances [1, 2]. The framework integrates the GridLAB-D agent-based distribution system simulation platform [3] with a human-behavior based residential load synthesis framework within a co-simulation environment. The tool is used to demonstrate how behavior models can be used to study the effects of the human factor in distribution level load management. Real-time (RT) demand response implemented with a residential load aggregator and a utility-scale battery controller are also demonstrated in this work. The real-time demand response technique described here employs device-level demand response in which customers grant utilities the ability to remotely adjust their energy consump-

tion by modifying the settings on internet-enabled devices such as smart thermostats and other similar control switches. However, in traditional demand response the devices are set on a static schedule with settings intended to lower consumption over the duration of a peak event. Here, demand response is used to control the activity of a feet of smart devices in real time to shape their aggregated power consumption in respond to some external control signal. In a different application, the residential load generator was incorporated into a real-time microgrid control system simulation as a proxy for a 200-home residential community with real-time demand response capabilities enabled by the residential load aggregator. The contributions by the author related to the implementation of the models within the real-time simulation are presented in this thesis as well.

1.2 Description of co-simulation framework

The use of DERs, and the use of electricity in general, is tightly related to how humans consume electricity across all application levels. As the technology to generate power from various distributed resources continues to improve, its economical appeal also continues to increase among customers in the utility, commercial, and residential sectors. A well-coordinated deployment of DERs coupled to the development of management systems will be crucial to maximize the benefits that these new technologies can provide. However, new DER management systems must be tested and analyzed to examine potential negative effects that they might impose on the grid. These test must be in such a way that avoids damaging expensive equipment or risks the quality of service to customers. The strong relationship between electric distribution systems and human behavior reveals the necessity for a simulation tool to conduct research that considers the effects of DER penetration on the electrical infrastructure, while also considering the critical role that human power consump-

tion and behavior plays within these systems. Part of the main focus of this thesis is a co-simulation framework that combines human behavior models with simulations of power distribution systems with the purpose of analyzing how human activities impact power management in distribution level systems. While models exist that simulate both of these domains individually, a combined model did not exist before the co-simulation tool presented here was published in. [1, 2]. Additionally, this framework is intended to serve as a tool in the development of new technology that will make the power grid more efficient, reliable, but most importantly, more resilient. Besides considering the effects of human behavior on the grid, it makes it possible to examine the performance of new control systems, and to analyze the effects imposed on electric infrastructure that would normally be neglected without a power flow simulator. For example, it provides the ability to precisely locate sections of a distribution feeder at risk of violating industry standards or experiencing faults, and to test potential solutions based on the insight gained from the results.

The co-simulation framework uses a residential load generator to simulate residential power consumption based on statistical representations of human interactions with household appliances. Residential loads are then passed to the GridLAB-D distribution system environment where power flow calculations are performed. Data transfer and time synchronization between co-simulators is accomplished using two different methods. The first employs the GridLAB-D *Matlab Link* directive function which uses a Matlab link control file to establish a connection to an external Matlab instance [4], allowing GridLAB-D to execute specific Matlab commands before main object and power flow calculations are performed during each time-step. The second method is an enhanced version of the first, taking advantage of FNCS (Framework for Networked Co-Simulation), a companion software package developed at Pacific Northwest National Laboratory (PNNL) that allows GridLAB-D to interact with other software as co-simulators [5]. For both methods, a typical scenario is examined first to confirm the co-simulations represent real distribution feeder behavior.

The scenario consists of a distribution feeder that hosts a utility-scale PV array and battery system, a residential community, and critical loads comprised of a hospital, supermarket, and water distribution infrastructure. Once the realistic behavior of the co-simulations is established, they are used to examine scenarios that consider the implementation of the residential load aggregator for demand response, and management of resources using the battery controller to manage the load on the feeder.

1.3 Description of real-time simulation

The increasing penetration of DERs has enabled power distribution systems to develop the ability to island sections, even entire feeders, within their system for microgrid operation. Microgrids host enough local power generating resources to support the power consumption of its customers, giving them the ability to disconnect from the main grid and operate without its support. While many microgrid systems exist today all over the world, the development of microgrid control systems is not mature. Most of the current controllers handle islanding requirements and provide management of generating resources and load balancing for resiliency and back-up purposes. However, the concept of the microgrid is changing to an advanced concept that includes all the essential elements of a large-scale grid, such as the ability to balance electrical demand with resources, maintain power quality, schedule the dispatch of resources, and preserve grid reliability in terms of adequacy and security. Therefore, modern microgrids require controllers with the capability to perform these functions [6].

A recent study conducted under a research project funded by the Mitsubishi Research Institute (MRI) with collaborative efforts from the University of New Mexico (UNM) and Sandia National Laboratories, investigates the application of a modern

microgrid control system in the Mesa del Sol community located in Albuquerque, NM. Mesa del Sol is the home of multiple DERs such as a utility-scale PV and battery storage system, the commercial-scale NEDO Microgrid that consists of a PV array, battery storage, fuel cell, and other DERs. The NEDO Microgrid is the legacy of a NEDO Microgrid Demonstration Project carried out by Japan's NEDO (New Energy and Industrial Technology Development Organization), the US Department of Energy, UNM, the local utility, several national laboratories, and many additional entities. After its completion, the microgrid was transferred to UNM and is now operated as a research testbed. The microgrid control system study considers real-life DERs currently available in the Mesa del Sol area as Hardware-in-the-Loop (HIL) components, as well as virtual models representing future resources such as residential demand response and electric vehicle (EV) charging stations.

In this case, the residential load generator was integrated into a real-time, hardware-in-loop, simulation as a virtual proxy for the existing Mesa del Sol 200-home residential community. In addition, the use of the residential load aggregator enables demand response capabilities for the virtual community. A virtual version of the Mesa del Sol distribution feeder was also used within the Opal RT real-time simulator to perform real-time, high resolution power flow simulation of the system. The microgrid control system links real and virtual DERs and uses a virtual power plant (VPP) approach, paired with a model predictive controller (MPC), for optimized management of resources. The VPP generates stochastically optimized day-ahead DER schedules based on load predictions and other parameters, while the MPC provides real-time dispatch of resources by deviating from the VPP schedules for an optimized use of resources based on current conditions [7]. However, an in-depth consideration of this microgrid control system falls beyond the scope of this thesis, and therefore, only the implementation and results related to the residential load generator and load aggregator will be discussed here. These include loading the models into their system platforms, requirements for real-time execution, communi-

cation between system components, and results of the simulated community as well as its response to demand response directives.

The real-time simulation platform showcased in this work demonstrates the ability to incorporate existing energy resources within a partially virtual environment that fosters the development of future infrastructure essential for modern microgrids. Most importantly, this platform enables the ability to test new microgrid control algorithms without risking equipment damage or quality of service to the customer. At the same time, it also provides the ability to gain valuable insight related to the introduction of new elements and their impact on existing electric infrastructure. Future use of this tool will help develop technology intended to make microgrid operations more effective, and to improve grid resilience by providing critical services in the event of severe disruptive events.

Chapter 2

The Models

2.1 Circuit 16 feeder model

Circuit 16, one of the feeders in the Los Alamos, NM (Fig. 1) power distribution network, was chosen for this study due to its utility-scale DERs and the diversity of its loads. It hosts a 1 MW PV array and a substation-sited battery energy storage system (BESS) composed of a 1 MW / 7.2 MWh NaS battery and a 0.8 MW / 1.2 MWh lead-acid battery. In addition to the utility-scale DERs, it serves a wide variety of loads including residential, commercial, and critical loads such as water pumps that contribute to the water distribution system of Los Alamos. Circuit 16 provides power to a total 1600 residential customers, while commercial consumers include two grade-level schools, a sports complex, and a golf course.

2.1.1 From WindMil to GridLAB-D

The Los Alamos Department of Public Utilities (LADPU), the local Los Alamos municipal utility, maintains a high fidelity model of the local power distribution system.

Fig. 1: The Circuit 16 feeder, outlined in orange, serves the Northeast area of Los Alamos, NM, covering mainly the Barranca and North Mesas. The battery storage and PV array are located at the local landfill and they interconnect to Circuit 16 at the substation, outlined in blue. Customers served by Circuit 16 include a 1600-home smart community, apartment complex, middle & elementary schools, a sports complex, a golf course, water distribution infrastructure, among others.

The model is based on Geographic Information System (GIS) data and accurate equipment specifications obtained from assets in the field and modeled in WindMil, a power utility software designed to model and analyze electric distribution networks for system planning and operations [8]. A total of 594 loads are defined in the model, representing the transformer-level power consumption of over 1600 customers served by this feeder. The reactance and resistance ratio, R/X, defined for the distribution lines within the model range from 0.025 to 2.36 and they distribute power to end users via overhead and underground lines with three-phase, bi-phase, and single-phase configurations depending on the type of load served.

The WindMil Circuit 16 model, shown in Fig. 2, was translated to a GridLAB-D

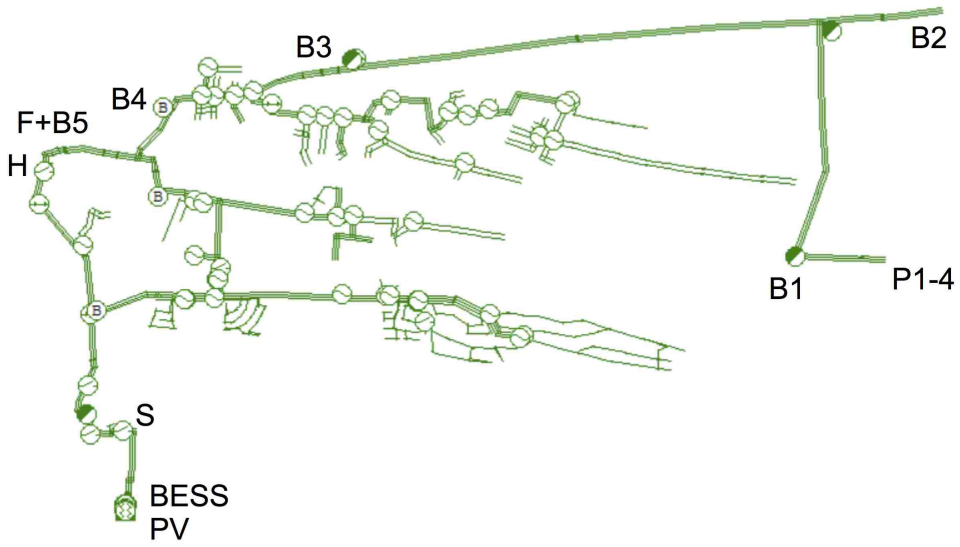

Fig. 2: Configuration of the Los Alamos Circuit 16 model, operated by the Los Alamos Department of Public Utilities. Note the presence of water supply well pumps (P1-P4), several pressure booster stations (B1-B5) and a water filter station (F). A 1MW PV array (PV), and 2 MW / 8 MWh battery storage (BESS) are located at the feeder head. Hospital (H) and supermarket (S) loads are also shown. Hundreds of residential loads are distributed at the feeder end nodes.

model using the export function within WindMil, allowing to translate the model from one platform to the other. Before exporting, Circuit 16 was isolated by suppressing the other feeders within the Los Alamos power distribution network model, ensuring that only the feeder of interest was translated to the GridLAB-D platform. However, this led to a few errors in the GridLAB-D model. For example, a *switch* object and its child *node* were created with mismatching phases. This was solved by defining the objects with matching phases. In addition, switches located where Circuit 16 attaches to other feeders were generated in a closed state leading to nodes that were not created during the translation process because they were defined within the suppressed feeders. Once these objects were identified and removed the GridLAB-D Circuit 16 model executed without errors.

2.1.2 Model validation

The Circuit 16 model is defined with a constant voltage of 66,395 V at the source node and stepped down via a transformer to 7,967 V at the head of the feeder. A snapshot of the voltage profile across the GridLAB-D Circuit 16 model is illustrated in Fig. 3. The graph on the left shows the voltage profile when serving only static loads. To assess how PV penetration affects the feeder's voltage profile, a solar generation resource was introduced near the end of a branch (at location B3 as shown in Fig. 2) to represent a 1 MW PV array. As expected, the PV penetration caused the voltage profile across the feeder to rise for all phases, but is especially noticeable on phase A. Further inspection of Fig. 3 also shows that the voltage increase is highest at the injection point. The case with no PV experienced minimum and maximum voltages of 99.25% and 102.37% of the nominal voltage, respectively, while the case with PV experienced a minimum voltage of 100.38% and a maximum of 102.86% of nominal. The fact that the voltage experienced in both cases is within ±5% of the nominal voltage confirms the Circuit 16 model meets ANSI C84.1 voltage range specifications, which states that the voltage throughout a distribution feeder must stay within ±5%

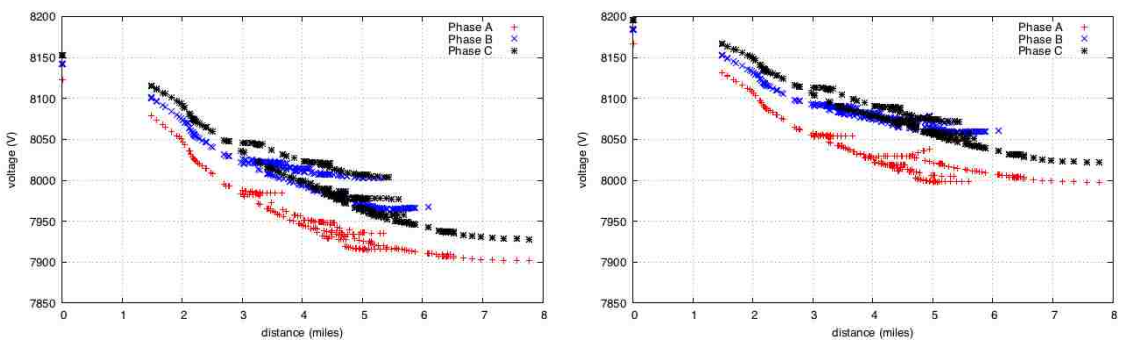

Fig. 3: Voltage on each phase as a function of distance from the head of the feeder. On the left is the case of power injected at the feeder head only. Power injected at feeder head and at the end of a branch by a *solar* object (at B3 for testing purposes, especially noticeable on phase A) is on the right.

of the nominal voltage [9] specified for the feeder.

2.1.3 Critical loads

The hypothesis of this work is that during emergency situations, a functioning critical infrastructure could be accomplished using distribution level DERs in combination with a distribution management system. As a result, this could reduce the stress imposed on transportation systems and emergency infrastructure during evacuation procedures by enabling a 'shelter in place' strategy. Since water, food, and medical care are basic needs for a community they are also considered in this thesis in an effort to replicate realistic loads on the feeder. Therefore, critical loads are considered to be a hospital, supermarket, and water distribution infrastructure.

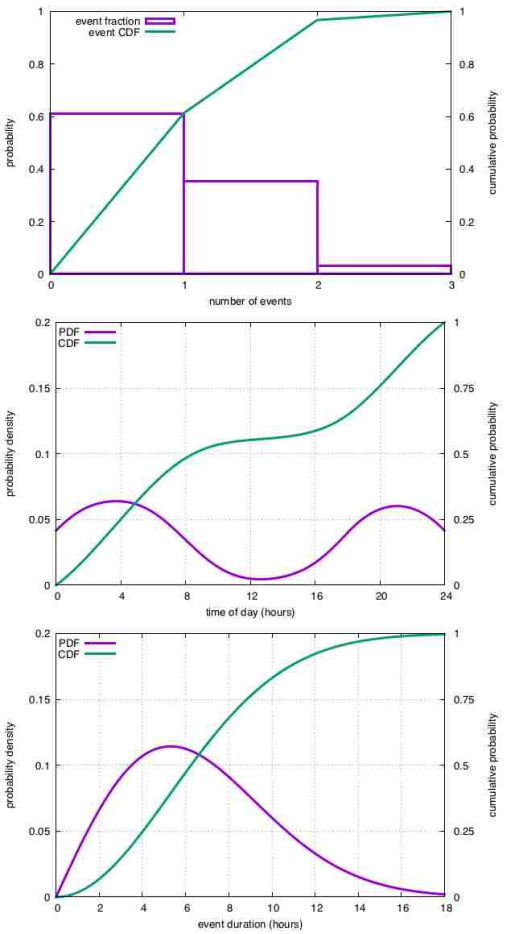

Fig. 4: Stochastic characterization of a water well pump in Los Alamos water system.

The water system load was developed using five months of operation data collected by LADPU from water pumps in the Los Alamos water system. These pumps contribute to the Los Alamos water supply directly and via a number of water storage tanks. A stochastic approach was used to generate models that simulate the water treatment load associated with Circuit 16. These models are characterized by probability density functions (PDF) correspond-

ing to the number of events per day, the duration of the events, and the start-time of the events experienced by each pump. The PDFs for the number of events per day are described by discrete probabilities, while load durations and event start-times are characterized by Rayleigh distributions and Gaussian Kernel Density Estimations, respectively. As an example, the probability density functions for a water well pump are illustrated in Fig. 4. In total, there are four water well pumps and five pressure booster stations included in the Circuit 16 model. The location of the water pumps on the feeder is indicated in Fig. 2, P1-4 and B1-5. Typical pump operation is shown in Fig. 5, with normal operation of the water pumps at night, when demand is low. The water is pumped from the wells to storage tanks distributed throughout the community, and dispensed during the day. High-resolution data for the hospital and supermarket are not available. Therefore, data for a hospital and a supermarket in Santa Fe, NM, a location with very similar climatic conditions, were used instead. The data correspond to the power consumption of these buildings during July 8,

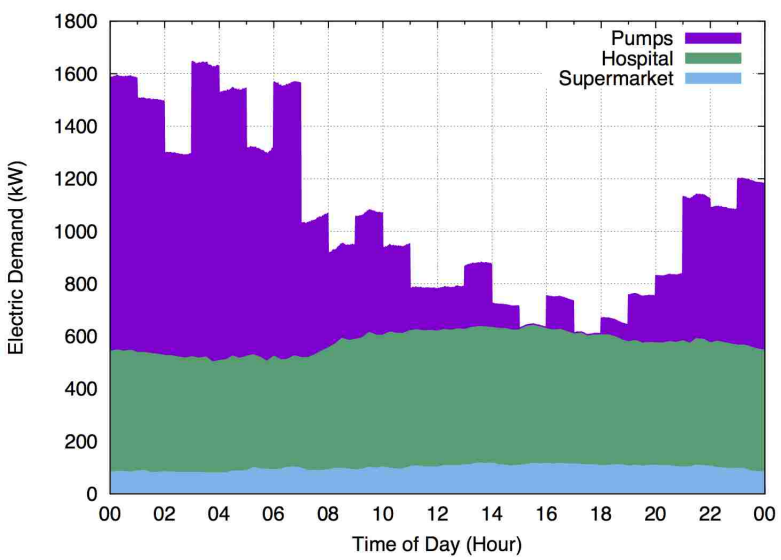

Fig. 5: Power demand of the critical loads for the Circuit 16 feeder, for July 8, 2016. Note that these loads, while supported by Circuit 16, are shared by the entire Los Alamos community.

2016. The electric demand profiles of these critical loads are also shown in Fig. 5.

2.2 Residential load generator & aggregator

The residential load generator & aggregator softwares were developed as part of a mutual research project shared with this thesis. The contributions made by the author to these models led to the ability to integrate them into multi-system simulations, operation in real-time, and enhanced input and output data, as well as their ability to communicate with other system components. The residential load generator and aggregator are introduced and described next, with a focus on the details related to this work.

2.2.1 Residential load generator

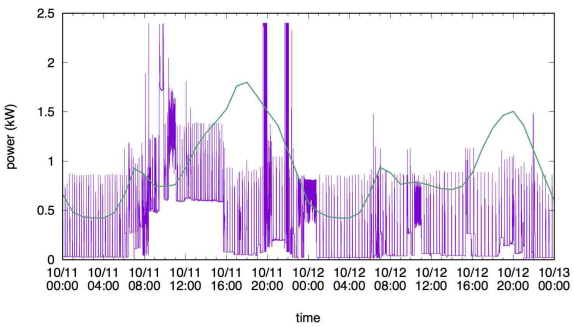

Fig. 6: Measured residential load, at one-minute resolution, compared to hourly building load simulation data.

To study the effects of the human factor on a distribution feeder, it is necessary to implement human behavior models that represent the interactions of people with the household appliances in their homes. The results from an hourly load simulation obtained using the framework described by Hendron and Engerbrecht [11] is compared to measured electric load data at one-minute resolution in Fig. 6. It is evident that hourly resolution data neglects considerable detail from the power consumption of a building. Hence, it is not reasonable to use hourly data in situations with dispatchable DERs

as it is often done in distribution simulations, given that dispatch of resources for power balancing must occur at least at a one-second resolution. Also, using data averaged temporally or topologically may lead to unrealistically smooth power flow and corresponding power quality metrics. In addition, local bottlenecks related to electricity infrastructure may not be revealed. This is especially true with high penetration of DERs, including PV generation, demand response, and storage.

Residential loads are associated with the use of household appliances that consume power ranging from a few Watts to a few kW. Here, a statistical approach is used to represent the interactions of human users with appliances and to account for the details of their operation, including the physical response of the house thermal management system. The use of a specific appliance is characterized by the number of events, the start-time and the duration, each described by a probability density function drawn from customer usage patterns. Statistical characterization of each appliance is also associated with a demographic cluster to account for different use of appliances by different people, namely working singles, working couples, families, and retired people. This is similar to the method described in the work of Fischer et al. [12]. The appliance package simulated for each home is composed of an electric range, refrigerator, lighting, HVAC, water heater, and a clothes dryer. The HVAC, water heater, and refrigerator models also contain thermodynamic systems, while the dryers and lights are considered ON/OFF loads, and electric ranges use a Markov-chain approach. An example of statistics for electric range operation is shown in Fig. 7. In this case, start-times correspond to meal times with peaks that indicate breakfast, lunch, and dinner. The magnitude of the peaks is associated with the different lifestyles in each demographic cluster. For instance, strong breakfast, lunch, and dinner peaks are specified for families, working singles and couples due to routine-driven lifestyles, while retired people usually enjoy unconstrained schedules. Durations PDFs are shorter for families during weekdays and longer for other clusters due to limited time to prepare meals. Longer durations are assumed uniformly

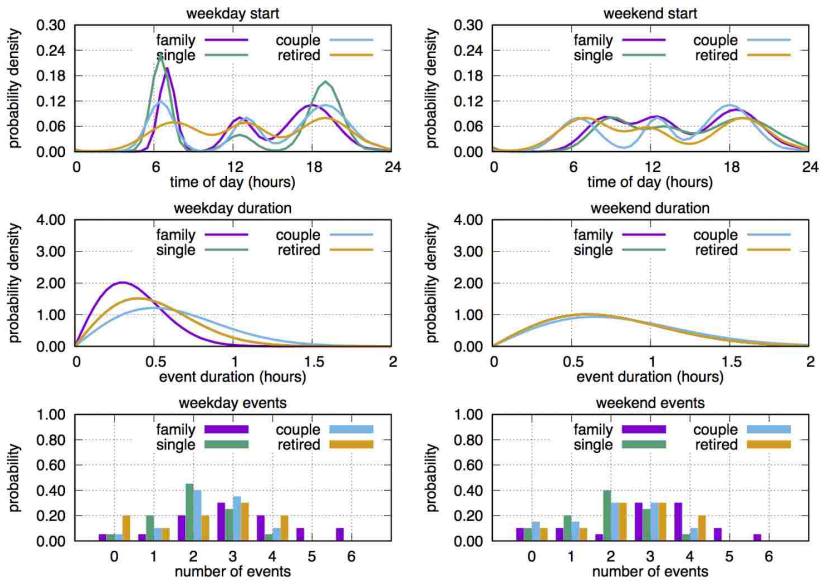

Fig. 7: Probabilistic characterization of the use of an electric range, in terms of probability density function for start-time, probability density function for duration, and discrete probability for number of use events. Note that PDFs and event probabilities are specified for four demographic categories.

for all categories during the weekend. The number of events is relatively high for families during weekdays, but retired people cook the most. However, the number of cooking events increases significantly for families on weekends due to a more people eating at different times of the day.

The interaction between the user and the physical nature of the service provided by the appliance is specific to the appliance. For instance, the interaction of the user with the air conditioning system is described by the temperature setpoints and deadbands set by the user on the system's thermostat. The following equations describe the interaction between the load and the physical system:

$$M_S \frac{dT}{dt} = \dot{Q}_L - \dot{Q}_R, \tag{1}$$

$$\dot{Q}_L = K_1[T_a - T(t)], \tag{2}$$

$$\dot{Q}_R = \Lambda \times COP \times P_{AC}, \tag{3}$$

where M_S is the thermal capacity of the conditioned space, $T(t)$ is the space temperature, T_a is the ambient temperature, t is time, Q_L and Q_R are the thermal losses from the structure and thermal gain from the air conditioner or heat pump, respectively, and K_1 is a constant related to the building's thermal insulation. Also, the coefficient of performance is denoted by COP, the power of the compressor is P_{AC}, and Λ is a state function that indicates whether the air conditioner or heat pump is ON or OFF. In its basic form, the air conditioning or heat pump is controlled by the switching logic

$$\text{if } T(t) < T_L \text{ then } \Lambda = 0, \tag{4}$$

$$\text{if } T(t) > T_U \text{ then } \Lambda = 1, \tag{5}$$

where T_L and T_U are the lower and upper deadbands for the temperature control. In addition to the air conditioner system, the loads of a refrigerator, a heat-pump type water heater, a clothes dryer, an electric cooking range and lighting are also considered. A similar thermal model to the air conditioner is used to describe the operation of the refrigerator. However, the user interactions with the refrigerator

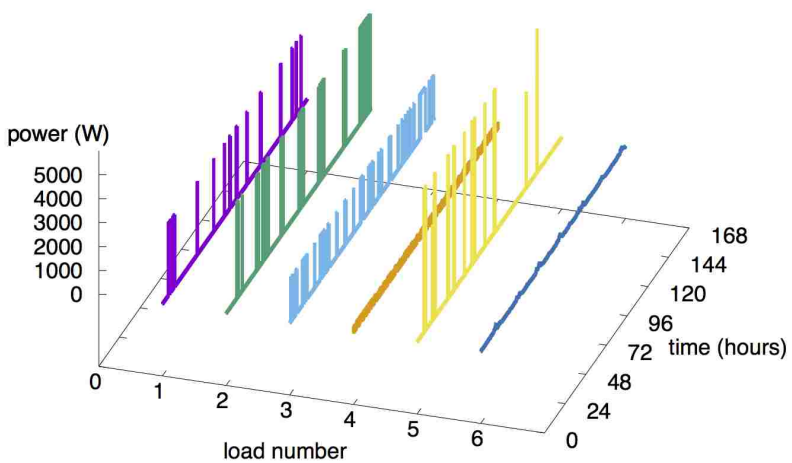

Fig. 8: Typical electric loads by appliance for a single meter, for a period of one week. The loads are clothes dryer (1), air conditioner (2), domestic hot water (3), refrigerator (4), cooking range (5) and lights (6).

are described as a function of open-door events instead of temperature setpoints. Similarly, the water heater is a thermal load with human interactions modeled as a function of water draws describing activities such as showers and dishwashing. The clothes dryer and the lights are modeled simply by their start-time and duration. If the event is active, a load is specified, otherwise there is no load. To simulate multiple lights in a house, multiple overlapping events are possible for lighting, while only one dryer event can be active at any given time with no overlap. In the case of the lights, the power of the LED light bulb is selected from a list of possible values. Furthermore, the Markov chain approach to simulate the lighting load used by Widén et al. [13], is employed here to model the power consumption of the electric cooking range. The typical power consumption of the household appliances considered here are shown in Fig. 8. Note that some loads, such as the clothes dryer, are high-power and relatively infrequent, while others are low-power but very frequent, such as the refrigerator. As a result, the total energy consumption by all the appliances is comparable. In addition, while the use of some appliances such as the refrigerator compressor are

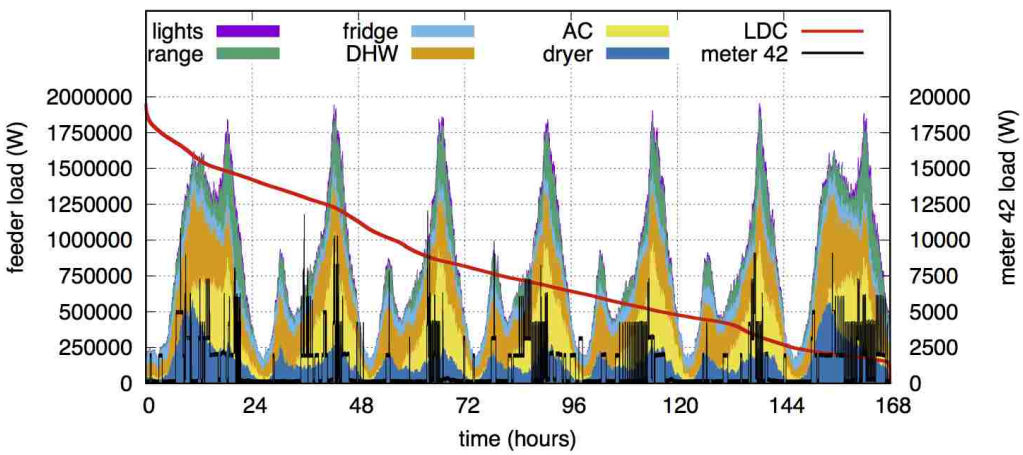

Fig. 9: Total feeder load resulting from 1000 houses, also indicating the aggregated contribution from each appliance. The total load for an individual meter (42) is also shown for comparison. The load duration curve (shown in red) indicates significant opportunity for shifting and deferring loads.

deferrable without detectable negative effects to the quality of service (QoS) to the customer, others such as the use of the cooking range cannot be deferred, but must still be characterized.

The total load of a residential community of 1000 houses is shown in Fig. 9 as an example. Also, the total load for a single home is shown in black for comparison with the total load of the community. It is evident that while the load of an individual meter displays significant variation in time with high variance, the total feeder load is relatively smooth and predictable. Also note that the power consumption of the individual home shown in this example is qualitatively very similar to the measured load shown in Fig. 6. Finally, the load duration curve of the 1000 house community indicates a number of interesting features of the load. First, a substantial peak load greater than 1.5MW occurs for less than approximately 10% of the time and could be easily removed by demand response in combination with storage. Second, the base load is only approximately 15% of the peak load. Third, the absence of a significant plateau indicates that there is opportunity to shift loads in time with appropriate control.

2.2.2 Residential load aggregator

In this study, a residential load aggregator manages the real-time demand response of the system. The appliances controlled by the aggregator are thermostatically controllable loads (TCL) such as the air conditioner (AC) and water-heating units in every home. For example, a $200m^2$ house has about 0.15kWh of stored energy per °C associated with the air mass alone, coupled with roughly 3kW of power controllable through a home thermostat. The load aggregator aggregates the total power consumption of the TCL units and calculates the capacity based on the state of the units in the system. The total power available to be decreased or increased through demand response depends on the number of units that are ON and OFF,

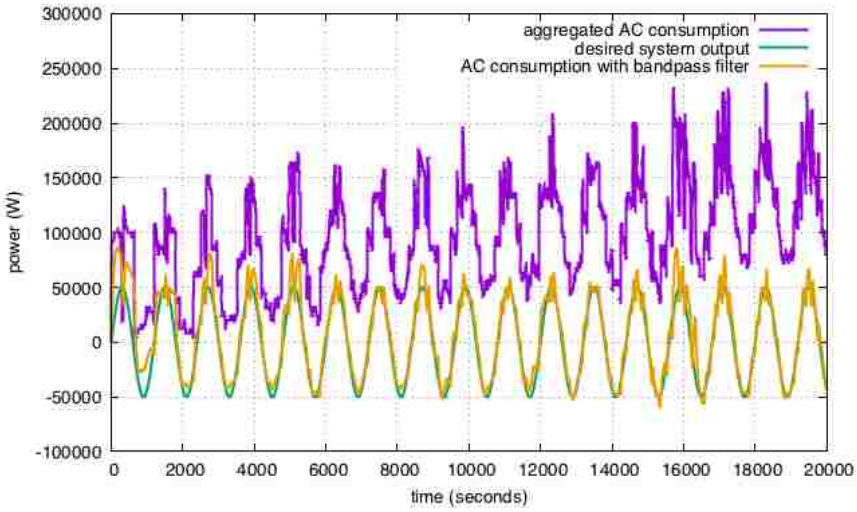

Fig. 10: Aggregated AC DR successfully following artificial sine wave signal. Total consumption of 200 AC units in in purple, filtered load and signal in yellow and green, respectively.

respectively. In this framework, instead of controlling each thermostat individually, the load aggregator broadcasts a single control signal to the entire fleet of thermostats in the system. The signal is in the form of a number that represents the probability of an available compressor switching from its current state. Furthermore, each "smart" thermostat decides to switch based on its current state and an internally generated random number. The broadcast signal is generated based on the error between the current power of the TCL system and a desired power setting that can vary dynamically. In other words, the system can be set to follow a static load shape on a schedule set by the utility to reduce the load on a feeder during peak hours, or it can be set to dynamically follow a changing signal such as solar irradiance. Fig. 10 shows the response of a TCL system composed of 200 thermostats following a demand response signal in the form of a sinusoid with a period of 20 minutes and amplitude of 50kW. This simulation test shows the system successfully following the artificial signal over a period of 20,000 seconds, or roughly 5.5 hours. The purple line shows the actual aggregated power consumption of the 200 AC units, which is

then passed through a bandpass recursive time series digital filter [14] to obtain the yellow curve for comparison with the sine wave. The filtering limits the response to only a specific band of frequencies for comparison with the demand response signal.

2.3 Battery controller

Utility-scale DERs such as batteries and PV arrays are vital resources in cases when the grid fails during emergency situation. They must have the ability to operate independently of the main grid in order to keep critical loads online. For example, in the case of a community such as Los Alamos where utility-scale DERs are available, batteries must be able to charge strictly from power generated by the PV array when power from the grid is not available. Also, power dispatch should be in such a way that optimizes the use of resources regardless of the state of the grid. The battery controller described here was designed to maintain the load at a point of interest, such as a transformer, at or below a predefined target power value by managing the power dispatched from available resources within the GridLAB-D powerflow simulation environment. It also enables the battery to charge directly from the power generated by an accompanying solar array. This battery controller was specifically created for this study since these capabilities were not available in GridLAB-D. Similar *load following* control schemes exist today in which batteries are set to charge or discharge at a constant rate when the load at the point of interest reaches certain thresholds. However, none combine the ability to charge a battery strictly from an accompanying solar array, manage the dispatch of available resources, and vary the discharge rate of the batteries according to the power needed to reach and remain at a target value.

The battery controller algorithm for this study was internally imbedded within the source code of the *inverter* object in GridLAB-D to which both the *solar* and *battery*

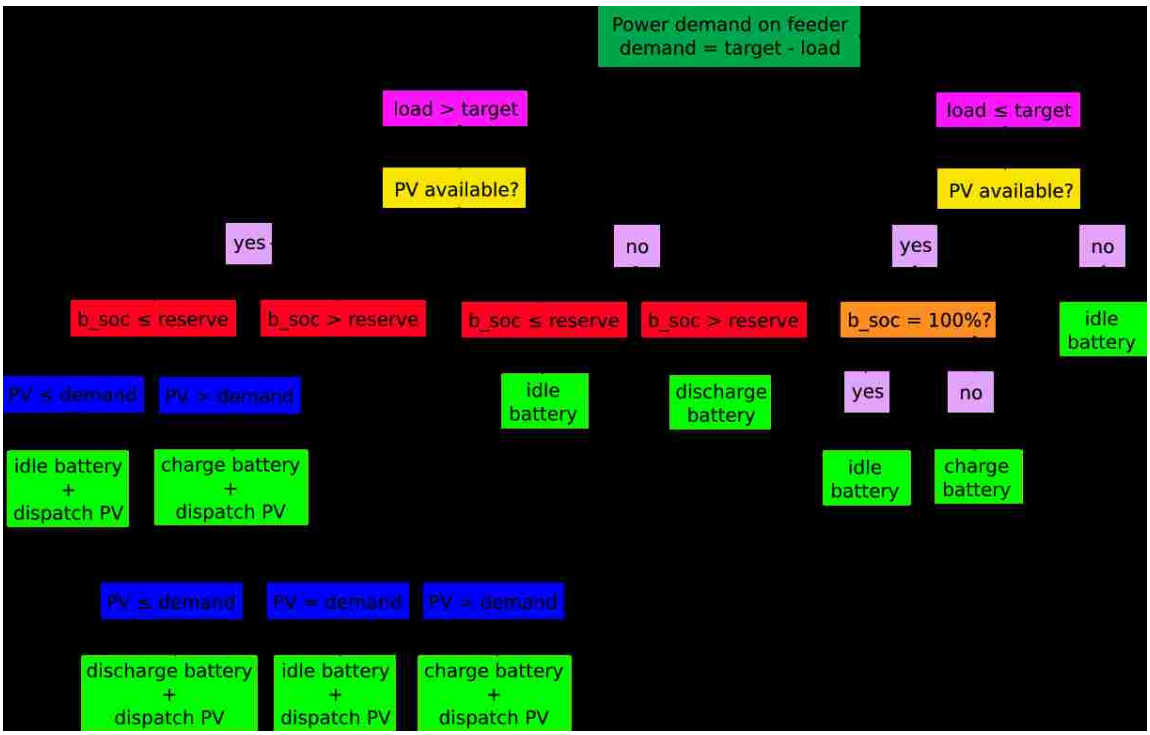

Fig. 11: Logic flowchart for battery control used to maintain the load at a point of interest on a feeder at or below a predefined target power value. The feedback control loop minimizes the difference between the current load at the point of interest and the target load. During the day, the controller dispatches PV and the battery if necessary; PV surplus is used to charge the battery. Only the battery is dispatched when necessary during the night, since PV power in not available.

objects connect to. The control algorithm uses the GridLAB-D simulation time handling process to create a feedback loop used to minimize the error between the target power value and the load on the point of interest. This method leads to the controller sensing the magnitude of the load as it was one second in the past, resulting in a one-second lag in controller action. However, this lag is negligible at the distribution level due to the relatively smooth load experienced by a substation. The controller dispatches PV during the day while charging the battery from PV surplus, and discharged only if it is necessary. During the night, only the battery is dispatched when necessary since PV power is not available. At this point, the battery is only allowed

to charge from a PV array; however, the ability to also charge from the main grid is within the scope of the future work discussed in Section 5.3. The flowchart in Fig. 11 shows the logic behind the battery controller. As a test, a simulation was conducted to examine the controller's performance. The test model consists of a single home connected to a bus to which a 500kWh battery and a solar array are interconnected for testing purposes. The PV and battery are defined with a rated power of 20kW and 10kW respectively. Fig. 12 shows the controller successfully maintaining the load on the primary meter at the predefined target load of 5kW for a period of two days, mainly due to the oversized capacity of the resources. The spikes above and below the target load are due to the one-second lag of the system and the resolution of the data collected. In a future version of this controller, the lag will be eliminated by calculating the power demand on the point of interest and forcing GridLAB-D to iterate the power flow solution for the time-step in which controller action takes place.

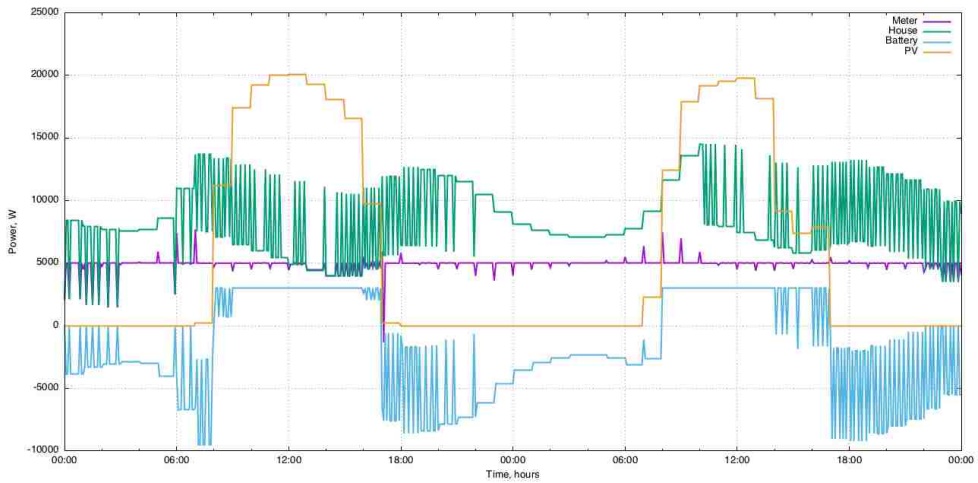

Fig. 12: Example of battery control performance on a test model exhibiting oversized DERs for testing purposes. The controller successfully holds the load on the primary meter at the target load of 5kW. The spikes above and below the target load are due to the one-second lag and resolution of the data.

Chapter 3

Simulation Setup

3.1 GridLAB-D co-simulation

A co-simulation that engages multiple incompatible models or software, requires an interface with the ability to handle communication and simulation-time synchronization between the softwares. Two co-simulation interface methods used to enable data transfer and synchronization between the *C++*-based GridLAB-D and the *FORTRAN*-native residential load generator and aggregator, are described in this section. The first method uses a GridLAB-D directive that enables communication with Matlab. The second method employs FNCS, a software especially designed to serve as the interface for co-simulations. Their setup and flow of information are described next.

3.1.1 Matlab as co-simulation interface

GridLAB-D's Matlab Link directive is used to integrate the residential load generator and aggregator with the GridLAB-D platform. This link directive enables Matlab

to execute specific commands on every time-step, before GridLAB-D performs main object and power flow calculations. This allows for the loads to be updated by the load generator in synchronicity with GridLAB-D. Generally, external loads are provided to GridLAB-D through *players* that read load-shapes from static files and apply the load to *load* objects. In the case of demand response, the load must evolve in parallel with the GrdiLAB-D simulation in response to controller directives issued by the aggregator. Here, Matlab fills the role of data communicator and synchronizer, while GridLAB-D acts as the master clock by utilizing its internal clock. Moreover, the load generator and aggregator are considered a single simulator within this setup, as illustrated in Fig. 13. That is, data is from the Matlab interface is received only by the aggregator, and the output data is sent back to the interface only by the load generator. Note that the communication between the aggregator and generator takes place independently, or outside, of the co-simulation interface. Instead, they exchange data via .csv files that also serve as communication signals.

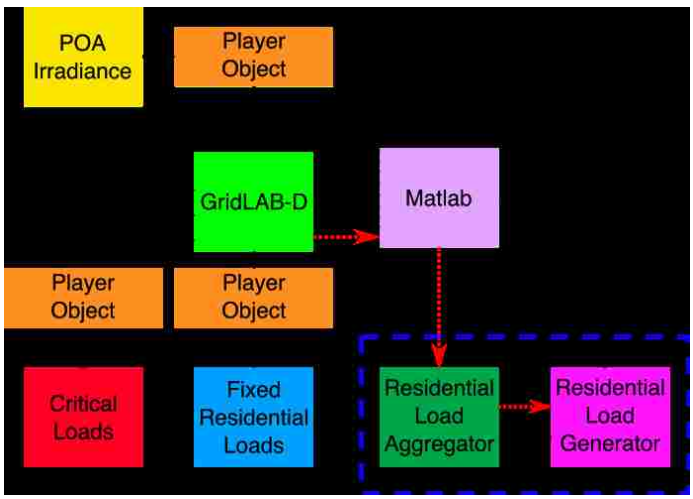

Fig. 13: Block diagram of information flow (solid black lines) between GridLAB-D and several input components, including non-controllable variable loads (via *players*) and controllable residential loads (via Matlab interface). Control signals are shown by dashed red lines. Note that the load generator and aggregator are considered a single simulator indicated by the blue dashed boxed.

In this exchange, the load generator sends the status of each unit participating in demand response and the aggregator provides the control signal for them to act upon. The co-simulation process is the following:

1) GridLAB-D initializes the simulation and prepares the environment for the first time-step, $t = 1$.

2) Matlab receives a signal from GridLAB-D to execute and begins by directing the load aggregator to run. It reads the status of the participating DR units and creates a demand response control signal (initial statuses are used on $t = 1$. Also, control signal is zero when demand response is disabled). The control signal is then stored in a .csv file for the load generator and its filename serves as a unique identifier.

3) The creation of the file containing the DR control signal also indicates the load generator to execute. It begins by registering the control signal and deleting the file. It continues by creating load values for the *load* objects defined in the Circuit 16 model and storing them in a .csv file with its unique identifier in the filename. A second .csv file is created containing status of DR units for the load aggregator to use in the next time-step. This is followed by a return signal sent to Matlab indicating that the loads have been created and stored successfully.

4) Matlab reads and imports the load values from the load file into GridLAB-D, updating the load values of the *load* objects.

5) Object and powerflow calculations are performed in the GridLAB-D environment for the current time-step t.

6) GridLAB-D advances in time to the next simulation time-step, $t = t + 1$.

7) Steps (2)-(6) are repeated until the simulation ends.

Here, the signals exchanged between GridLAB-D, Matlab, the load generator, and the aggregator are crucial for synchronization of the overall simulation. They ensure that steps (2)-(6) are performed only once per time-step and not every time GridLAB-D iterates when solving powerflow across the feeder. Once again, the load generator and aggregator are treated as a single simulator unit within this setup, as indicated by the blue dashed box in Fig. 13. To simulate a community with demand response capabilities within Circuit 16, 50 residential meters are simulated by the load generator and aggregated to 10 load values that represent the number of residential transformers in this community. To account for the 594 residential loads originally defined in the model, a static load-shape is used for the remaining 584 fixed residential loads that are defined within the GridLAB-D model. Scale factors are used to ensure that the maximum power consumption of each *load* object under the use of the load-shape, is comparable to their original constant power consumption defined in the model. Static load-shapes are also used for the load associated with critical infrastructure and solar irradiance, and are fed into GridLAB-D using conventional *player* objects.

3.1.2 FNCS as co-simulation interface

The FNCS co-simulation interface is an enhanced version of the Matlab interface as they were both built with the same purpose. The FNCS interface allows for a higher data transfer capacity and speed with stronger reliability and performance. In a similar way as in the Matlab interface setup, the GridLAB-D Circuit 16 model is integrated with the residential load generator and aggregator within a co-simulation. Here, the FNCS software allows for a co-simulation in which residential loads are updated by the load generator in synchronicity with GridLAB-D and the load aggregator. Similar to the Matlab setup, loads must evolve in parallel with the GridLAB-D simulation in response to demand response directives issued and managed by the ag-

gregator. In this setup, the FNCS broker represents the central server to which all simulators connect to in order to synchronize in time and exchange messages with other simulators [15]. The broker acts as the master clock, simulation synchronizer, and data communicator between the three co-simulators. The simulation process is the following:

1) The co-simulators including GridLAB-D itself, are initialized and they prepare their environment for the first time-step, $t = 1$. This includes connecting with the FNCS broker and subscribing to publications that will be made by other co-simulators.

2) The FNCS broker receives the status of each simulator and sends the signal to execute for the current time-step only when all co-simulators are ready.

3) The simulators receive the signal from the broker and begin by requesting the published values to their subscriptions from the broker. This happens at the beginning of each time-step (initial conditions are used for $t = 1$).

4) The simulators execute for the current time-step, t, and publish their output to the FNCS broker, which then brokers the publications to the respective co-simulator subscription.

5) Steps (2)-(4) are repeated until the simulation ends.

It is important to note that, unlike in the Matlab interface setup, the load generator and aggregator are treated as individual simulators and communicate only through FNCS, as illustrated in Fig. 14. Furthermore, all simulators obtain values to their subscriptions at the beginning of each time-step as indicated by step 3. Here, only the solar irradiance and the critical loads associated with the hospital, supermarket, and water distribution are static load-shape fed into GridLAB-D by *player* objects,

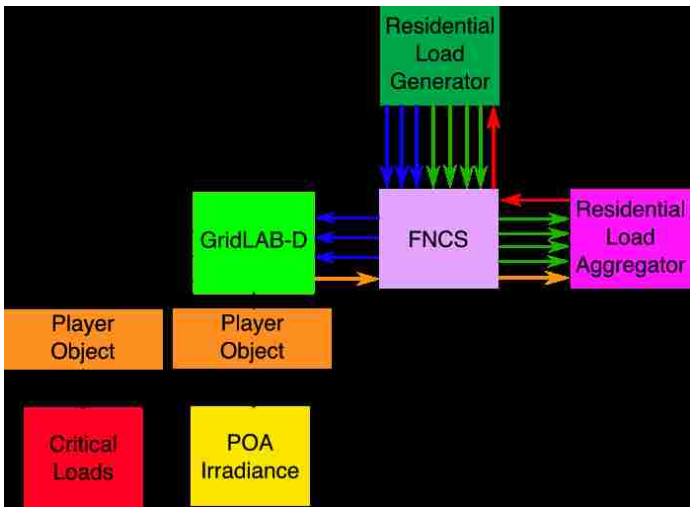

Fig. 14: Block diagram of information flow between FNCS and co-simulators, including GridLAB-D, residential load generator and aggregator. Critical loads and plane of array irradiance are read by GridLAB-D via *players*.

similar to the Matlab setup. Due to the higher data transfer capacity, the load generator simulates a total of 1600 homes aggregated to 594 values that represent the number of residential homes and transformers within the Circuit 16 feeder, respectively. The load generator publishes the 594 aggregated values for the GridLAB-D Circuit 16 model in addition to TCL status of each home for the aggregator. In return, the aggregator publishes the control signal for the load generator indicating the demand response directive. Finally, the Circuit 16 model publishes the solar irradiance used by the aggregator to generate the demand response signal. Overall, a total of 2,196 values are transferred between three co-simulators each simulated second.

3.2 Real-time simulation

The residential load generator and aggregator were integrated into the real-time simulation described in Section 1.3 as a virtual proxy of a demand response enabled community of 200 homes. To simulate the physical separation of the simulated residential infrastructure and the residential load aggregator in real life, the models are executed on individual Raspberry Pi 3 computers (Fig. 15). These are credit-card-sized single board computers that run on Linux operating systems and offer Internet connectivity [16]. The simulation employs a central database as the interface between real and virtual components responsible for real-time information used for operations optimization and control. Communication between the models and with the database is performed via the Internet to replicate communication of these devices in the field. Furthermore, the models interact with the database using a Python-based interface developed by Sandia National Laboratories, while communication between models takes place using the scp (Secure Copy) command line tool. This tool copies files across remote hosts using a secured encrypted connection over a network [17, 18]. Fig. 16 illustrates the simulation setup between the load generator, the aggregator, and the database, as well as the total flow of information that takes place on every time-step of the simulation.

Fig. 15: Raspberry Pi 3, credit-card-sized single board computers running Linux operating systems.

The simulation was set up to run in real-time with 1-second resolution in order to accurately represent the power consumption of the Mesa del Sol 200-home residential community, along with demand response controllability of TCL units described in Section 2.2.2. At the beginning of each time-step, the residential load genera-

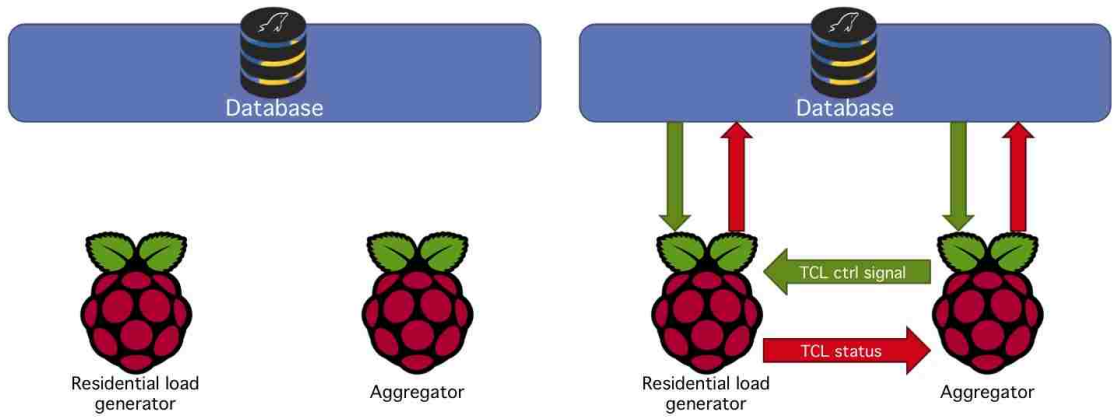

(a) Real-time simulation setup. (b) Real-time data transfer diagram.

Fig. 16: Real-time simulation setup showing the models and their interaction with the database. (a) The residential load generator and aggregator reside in individually independent Raspberry Pi 3s which communicate with the centralized database. (b) Input and output data flow is shown with green and red arrows, respectively.

tor and the aggregator request the current outside temperature and a MPC setting value from the database, respectively. Recall from Section 1.3, that the MPC provides real-time dispatch of resources by deviating from schedules provided by the VPP. Moreover, the load generator and the aggregator exchange TCL data critical for demand response. In this data exchange, the load generator receives a control signal from the aggregator as an input and returns the current statuses of the TCLs to the aggregator. The aggregator takes the TCL statuses and the MPC setting it receives from the database as inputs, and determines the control signal based on these values. Samples of the TCL status sent to the load aggregator are illustrated in Table 1, where the first column represents the meter number of each house, the second and third columns indicate whether the AC unit of the home is ON and if it is available to switch states (from ON to OFF, and vice versa), respectively. A value of 0 indicates that the AC is OFF and a value of 1 indicates it is ON. Similarly, 0 signifies that a given AC unit is not available to switch states, and the opposite is signified by a value of 1. The last three columns inform the aggregator of the overall

meter_id	ac_on	ac_switchable	ac_soc	lower_db	higher_db
1	0	1	0.288380	16.000000	26.000000
2	0	1	0.426664	16.000000	26.000000
3	0	1	0.339106	16.000000	26.000000
4	0	1	0.372218	16.000000	26.000000
5	0	0	0.033459	16.000000	26.000000
...	...	...	...	...	...
196	0	1	0.365487	21.742659	23.742659
197	0	1	0.219987	21.784323	23.784323
198	0	1	0.086933	21.990416	23.990416
199	1	0	0.245799	16.000000	26.000000
200	0	1	0.066387	16.000000	26.000000

Table 1: Sample of TCL status data passed from the residential load generator to the aggregator for 200 air conditioners. The data include the identification number for each meter, the current state of the unit (ac_on), availability to switch states (ac_switchable), the state of charge of the unit (ac_soc), and the lower and upper deadbands (lower_db & higher_db).

state of charge, as well as the lower and higher deadband setpoints of each unit.

The flow of input data at the beginning of each time-step is depicted in Fig. 17a and the data flow at the end of each execution between the load generator, the aggregator, and the database is illustrated in Fig. 17b. The load generator aggregates the power consumption of the 200 homes into 38 load values that correspond to the 38 residential transformers located in the Mesa del Sol community. The total power consumption of the community is also calculated and sent to the database along with the 38 transformer values. Each value is given a database load identification number (load_id) along with the timestamp of the current time-step of the simulation. Table 2 shows an example of this data in its final format before it is sent to the database. Note that the first row, load_id=2, corresponds to the total power consumption of the community and the rest correspond to the transformer values. Also, load_id=1 is associated with the load of an existing commercial building in the area. At the same time, the aggregator sends the TCL's current state of charge,

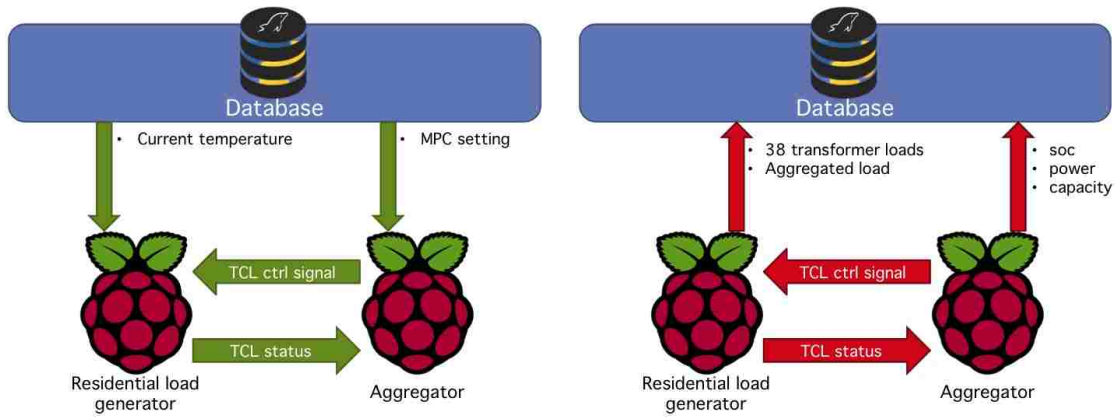

(a) Input. (b) Output.

Fig. 17: Real-time simulation input & output data flow. (a) The load generator and aggregator request data from the database as their input. They also exchange TCL data for DR. (b) The simulators send their output to the database at the end of each time-step.

power, and capacity to the database to be used by the MPC and VPP.

In order to make the models execute in real-time it was necessary to synchronize the simulation with the current time, while also ensuring that the models execute a time-step every second. Synchronizing with real-time requires each model to keep track of the current date, time (with a resolution of 1-second), and the current day of the week. The day of the week is important due to the difference in residential power consumption behavior during weekdays, when people are usually at work, then on weekends, when people are more likely to be at home. At the start of the simulation each model requests the clock value of its host machine which provides the current date and time. The models then take a difference of days between the current date and a reference date for which the weekday is known. Here, the reference date used is Sunday January 1st, 2017. With this information and the fact that there are seven days in a week allows the models to calculate the current day of the week and to represent it with an integer from 1-7, where 1 represents Sunday. Moreover, the time obtained from the clock query indicates the start hour, minute, and second of the

load_id	datetime	measured_value
2	2017-08-23 11:47:03	126269.070
3	2017-08-23 11:47:03	10334.182
4	2017-08-23 11:47:03	4540.000
5	2017-08-23 11:47:03	758.000
...	...	...
36	2017-08-23 11:47:03	532.000
37	2017-08-23 11:47:03	7290.000
38	2017-08-23 11:47:03	6600.182
39	2017-08-23 11:47:03	5342.182
40	2017-08-23 11:47:03	2350.182

Table 2: Sample data transferred to the database by the residential load generator, representing the total aggregated load (load_id=2), and followed by individual transformer loads, for use in the Opal RT real-time simulator. Note that load_id=1 is associated with the load of a building outside of the scope of this work.

simulation, as well as the number of time-steps left for the current day. In order to keep track of time and generate the correct timestamp, the clock value is requested from the host machine on every time-step. To ensure that times remain uniform across components, each machine obtains the global time from a network time protocol server via the internet.

Ensuring that the models execute every second required the development of a trigger that indicates to the models when to run. The Python-based trigger used here is independent of the simulation and its components, and runs quietly in the background keeping track of the computer's clock. It creates a trigger file that signals the models to execute each time the seconds digit of the clock changes. Each model is associated with an individual trigger that lives on the same machine and is responsible for signaling its corresponding model to execute for the current time-step. At the end of the execution process, the models delete the trigger file in sub-second time intervals. Note that the trigger file is created regardless of what is happening within the simulation; therefore, proper error checks and error-handling logic were necessary and are crucial to ensure that the simulation keeps running even if fatal errors are en-

countered. These error-handling procedures were designed with field deployment in mind. For example, the load aggregator would normally crash if it was signaled to execute a time-step, and the current TCL status data was not available. Instead, the models are instructed to reuse the data they had on the previous time-step if the new input data is missing or if was not updated in time.

Chapter 4

Simulation Results

4.1 GridLAB-D co-simulation

Several scenarios were examined using both of the GridLAB-D co-simulation interface setups described in Section 3.1. In both cases, a preliminary simulation was performed to confirm that the combined models represent real distribution feeder behavior. Once realistic behavior is established, scenarios are performed using the battery control and residential demand response controllability. All scenarios considered using such co-simulation interfaces are simulated to take place during July 8, 2016, the day corresponding to the load-shapes used for the critical loads. In addition, a hot summer day is ideal to stimulate the real-time demand response control of AC units and observe its performance.

4.1.1 Matlab interface

Three cases are considered using the Matlab interface. First, the typical 'fixed schedule' approach, in which the utility battery is charged at night and discharged during

peak hours, with no additional controls. In the second case, the battery controller from Section 2.3 is introduced. It uses load feedback from the substation transformer to dispatch the utility-scale battery in order to limit the peak load on the feeder. In the third case, an example of demand response is demonstrated by aggregation of household water heaters. The residential load is composed of a static load-shape fed into 584 *load* objects, and scaled such that the peak power consumption of each load is equal in magnitude to the constant power value with which they were originally defined in the model. In addition, 50 residential homes are aggregated to 10 load values that are passed to GridLAB-D for the total 594 *load* objects. They represent the number of homes and transformers located within a section of a simulated community with demand response capabilities. The solar irradiance and the critical loads associated with the hospital, supermarket, and water distribution are also static load-shape read by GridLAB-D using conventional *player* objects. Outside temperature is represented by a sinusoid, $T_{amb} = 28 - 8sin\frac{hour*\pi}{12}$, with characteristics similar to the natural temperature cycle experienced throughout the day. The

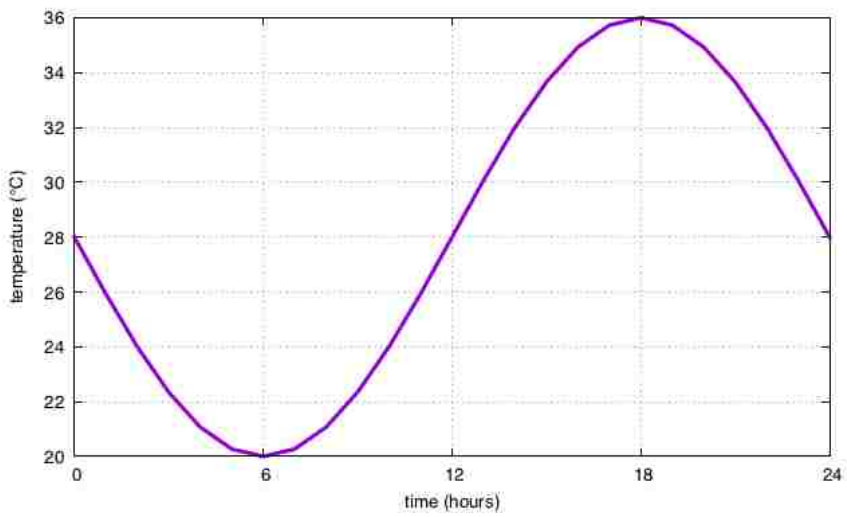

Fig. 18: Artificial outside temperature representing a hot summer day in Los Alamos, NM. Maximum artificial temperature is comparable to the maximum recorded temperature for Los Alamos. The temperature profile is represented by a sine wave.

temperature is set to range between 20-36°C to represent a maximum temperature similar to the highest recorded temperature in Los Alamos, NM of 35°C in 1998 (Fig. 18). Note that the purpose of the Matlab co-simulation setup was mainly to establish the ability to integrate models that cannot operate satisfactorily together (incompatible models) under typical computer conditions. The goal is to perform co-simulations that consider aspects from different domains, such as human behavior and power flow principles, in order to simulate accurate power distribution behavior and to ensure that the simulated infrastructure operates within physical limits.

A. Uncontrolled conditions

Power distribution feeders equipped with DERs such as PV and batteries usually discharge the batteries during peak load hours to reduce the stress on the transformer, while charging takes place at night when wholesale energy cost is low. This mode of battery operation is shown in Fig. 19. The typical battery schedule indicates the battery to charge from midnight until 07:00 at 1/3 of its power capacity ($\sim$ 540kW)

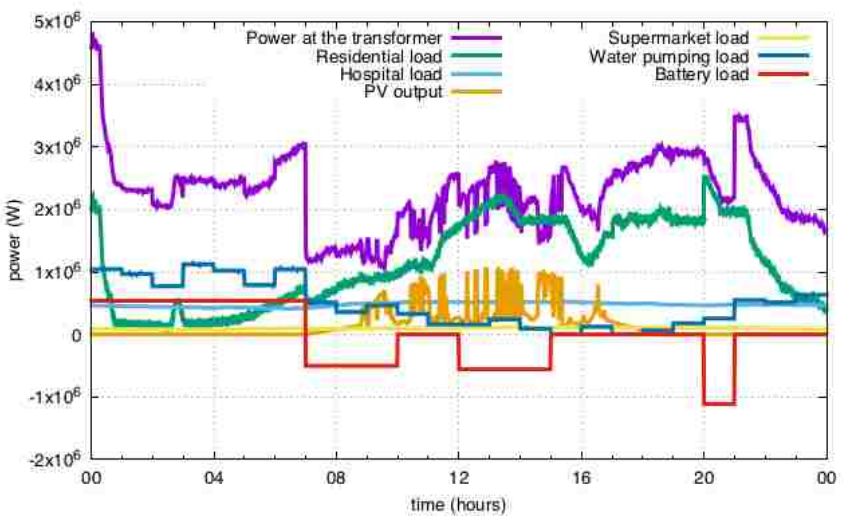

Fig. 19: History of net load at the feeder head, in comparison with critical loads, residential loads and battery charge / discharge on July 8, 2016.

using power from the grid. It then dispatches power at 1/4 capacity (500kW) from 07:00-10:00, in order to aid with the typical morning peak associated with consumption as people prepare for work; shower, breakfast, etc. The battery is set to idle for two hours before dispatching power again from 12:00-15:00 at about 556kW to assist with the typical lunch peak. This is followed by another five-hour idling period while most people are at work. Furthermore, another peak occurs around dinner time associated with people returning home from work, cooking dinner, doing laundry, etc. Thus, the battery is set to dispatch power from 20:00-21:00 at 56% capacity ($\sim$ 1.1MW) in order to lessen the load on the feeder associated with this behavior. It is also important to note from Fig. 19, the intermittency of the PV array associated with a relatively cloudy day. An intermittent irradiance profile was used in an effort to deviate from ideal solar conditions in order to observe interesting characteristics of the system that might not be easily apparent under ideal conditions.

A closer look at the history of the loads shown in Fig. 19, indicate that while the residential load consumption of 1600 houses is low during the night, it dominates during the day. Comparing the size of the residential load to the critical loads during the day confirms the potential for aggregated residential demand response as an effective tool to ensure that power is available to serve critical loads. Moreover, the loads associated with pumps used to distribute water to the community are high at night to take advantage of low energy costs and also to manage total feeder load effectively. While this is a good strategy for normal conditions, moving pump loads to the daytime, when PV power is available, may be a preferable strategy.

B. Control with battery

In some applications, more complex control algorithms are used instead of fixed schedules to maximize the benefits of a battery. A demonstration of this kind of operation is described in the work of Lavrova et al. [19, 20], in which a framework of

control techniques was used to optimize the benefits of battery in conjunction with PV at the distribution scale.

The performance of the GridLAB-D built-in battery controller described in Section 2.3, is demonstrated here. Results from the co-simulation illustrate the dispatch of PV and battery power managed by the inverter. Note that all remains identical to the previous case except for the operation of the battery. Here, the main goal of the controlled operation of the battery is to maintain the load at the substation transformer at or below 2MW for the duration of the day. The controller dispatches PV and if necessary the battery to ensure that the transformer load remains below the predefined threshold as shown in Fig. 20. The initial battery SOC (state of charge) is assumed to be 0.9 at midnight, resulting from control actions implemented the previous day. It is evident that the control is effective until shortly after noon, when battery capacity runs out after reaching its reserve SOC of 25%, in part due to lack of PV power owing to partly cloudy conditions. After that, loads are supported pri-

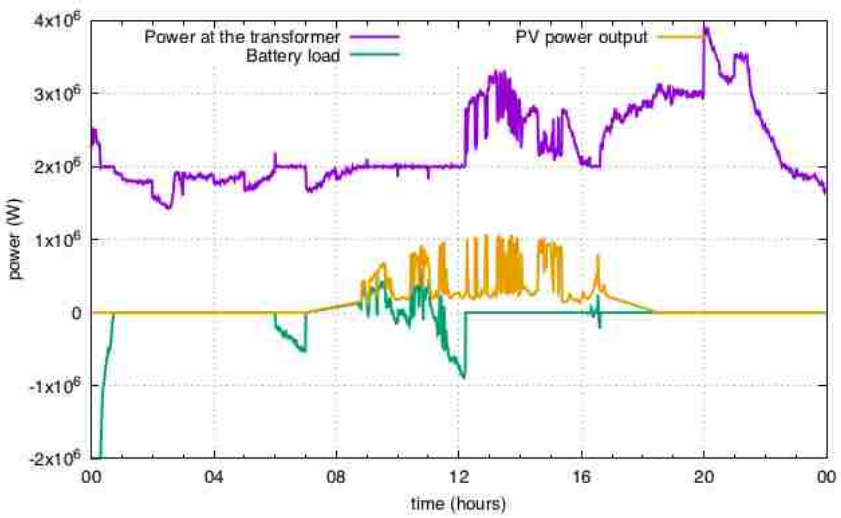

Fig. 20: Performance of battery control during the day of July 8, 2016. The battery control the load at the substation transformer effectively at the pre-set maximum value of 2MW until the minimum SOC of 0.25 is reached at around noon. After that, loads are met via the PV and the grid only.

marily by the grid, with partial PV support as dictated by sky conditions. While this is generally not a problem when the grid connection exists (unless maximum transformer load is violated), operation of critical infrastructure in emergency conditions could not be supported. The insight gained from these results highlight the value of this co-simulation framework as a tool to study the effects of more sophisticated control algorithms on the electrical infrastructure.

C. Control with DR

There are two possible methods to support critical infrastructure operations. One is the availability of stand-by fossil generation in combination with larger battery storage. The other is curtailment of non-critical loads. The former option is expensive and not likely to be implemented even if the value of resilience were built in to the infrastructure cost. The second option is more cost-effective, but would require the implementation of a control infrastructure to manage the partial operation of non-critical loads. While the design of such an infrastructure is outside the scope of this thesis, an example of how the simulation framework used here could assist in the design process is given below.

The residential load aggregator is introduced here to perform demand response on 50 residential homes by controlling their water heating units simulated within the load generator. The aggregated load of the 50 meters is shown in Fig. 21. The initial state of the heat pump water heaters is OFF, with an initial SOC normally distributed around 0.25 with a standard deviation of 0.05, and a SOC deadband of 0.9 to 1.0. The minimum deadtime conditions of compressor operation are met after 5 minutes, indicating the heaters to turn ON and begin the charging process. In the meantime, the GridLAB-D controller sends a signal to the load aggregator indicating to reduce residential water heating loads. However, this is prevented internally by the aggregator due to the SOC value being below the deadband minimum. At 30 minutes,

the internal deadband minimum is changed to 0.4, allowing the aggregator to turn OFF heaters that are ON according to the control signal. The aggregated control results in an immediate load reduction of 50kW within 1 minute, and a reduction of almost 100kW for the 50 houses within 10 minutes. With higher participation, this effect would be magnified accordingly. For example, a load reduction of about 500kW would be possible with demand response control of a 250-house community. In the case of an unexpected emergency, a utility could employ this technique to reduce the load on a feeder while a backup generator is initiated for support. Similar algorithms could be implemented for laundry loads and other non-essential appliances to reduce load during emergency conditions. An example of air conditioning control is discussed in Sections 4.1.2.B and 4.1.2.C.

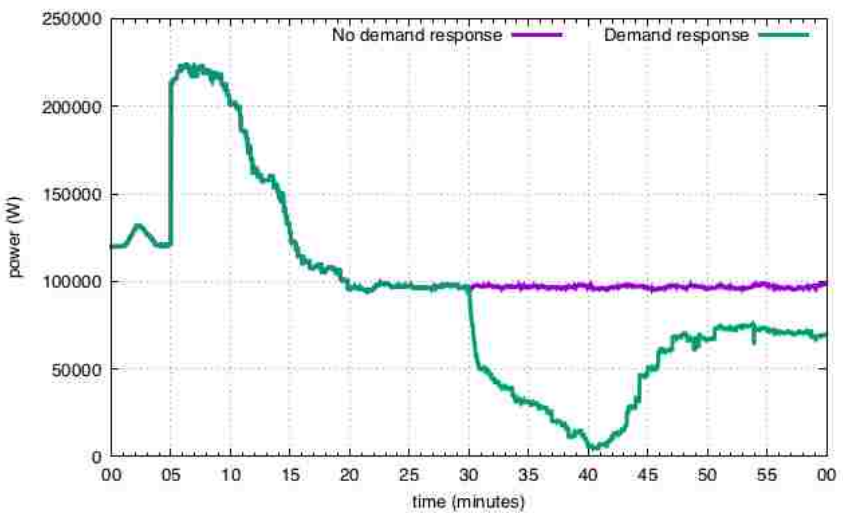

Fig. 21: Example of aggregated DR of 50 domestic hot water heaters. At 30 minutes, the lower deadband limit for the state of charge is changed from 0.9 to 0.4, allowing the external signal from the aggregator to turn off available heat pump compressors.

4.1.2 FNCS interface

Here the FNCS interface was used to examine three cases. The first consists of typical operating conditions used to establish realistic distribution feeder behavior of the system, similar to the uncontrolled case performed using the Matlab interface. In the second case, demand response is enabled using solar irradiance as a signal to control a fleet of air conditioner systems within 1600 homes. The third case demonstrates the performance of demand response coupled with the battery controller. The main difference between the Matlab interface and FNCS, is the amount of data shared by the simulators. In this case the load generator simulates a total of 1600 homes aggregated to 594 values that represent the number of residential homes and transformers within the Circuit 16 feeder, respectively. Also, demand response is applied to the entire residential community, as opposed to only a small section of a community as seen in Section 4.1.1.C. Furthermore, temperature data recorded on a mild summer day in Phoenix, AZ, shown in Fig. 22, was used here to induce stronger demand response effects. However, similar to the Matlab setup, solar irradiance and critical

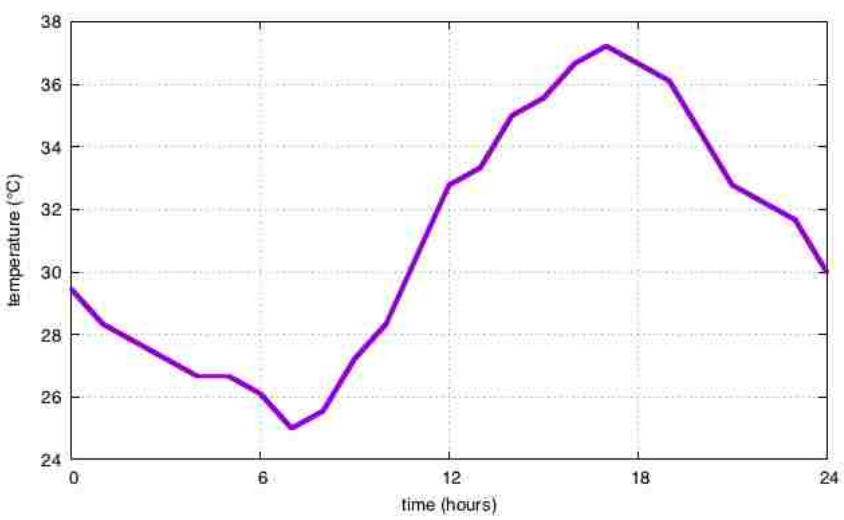

Fig. 22: Temperature recorded on a mild summer day in Phoenix, AZ. It is used here to intensify the effects of demand response control of 1600 AC units.

loads remain as static shapes.

A. Baseline

In order to establish a performance baseline for the system, the same typical mode of operation for distribution feeders described in Section 4.1.1.A is used here. Note that the battery is set on the same dispatch schedule discussed in Section 4.1.1.A, indicated by the red line in Fig. 23. Moreover, PV generation, as well as the loads associated with critical infrastructure such as the hospital, supermarket, and water pumps, remains the same as in the Matlab benchmark scenario. However, in this case the entire residential load is simulated by the load generator composed of 1600 homes. The residential load results in a typical profile with low consumption during the night, and high power consumption during the day.

To examine the effects of DER penetration to the feeder, a voltage profile was generated during maximum PV generation and battery dispatch totaling roughly 1.6MW

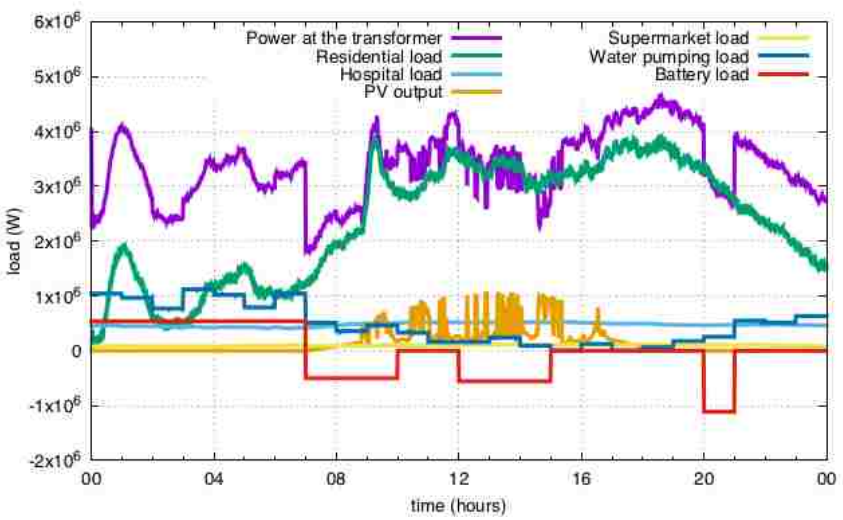

Fig. 23: History of net load at the feeder head, in comparison with critical loads, residential loads and battery charge/discharge on July 8, 2016 using the FNCS interface.

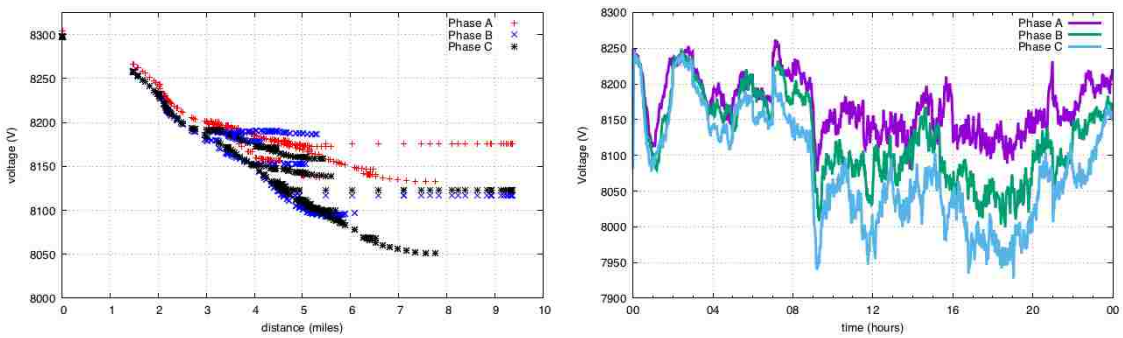

(a) Node voltage at 12:55.

(b) Voltage history of a branch.

Fig. 24: Voltage behavior on Circuit 16 during baseline case. (a) Profile of node voltage as a function of distance from the source node experienced by Circuit 16 during high PV penetration and uncontrolled conditions, 12:55. (b) Voltage history of a node located midway on the branch below location B3.

provided by DERs. Fig. 24a shows the voltage at every node as a function of their distance from the source node at 12:55. The maximum and minimum voltage values are 104.2% and 101.1% of the nominal voltage of 7967.43V, respectively. While these values are within the ANSI C84.1 voltage range specifications, the maximum voltage value is within 0.8% of violating the upper voltage limit. It is evident from Fig. 24a, that this is due to DER dispatch since the maximum voltage is experienced at the substation where DERs reside. This suggests that DER penetration is approaching its limit due to high PV generation. At this point, it would be necessary to lower the support from the battery to avoid any violations. Ideally, the battery would be set to charge in order to take advantage of the power generated by the PV array. Furthermore, this points out the importance of more sophisticated controls that automate battery activity in order to replace fixed schedules and limit human interaction, while improving performance and anticipating potential violations to industry specifications.

Additionally, to examine how other parts of the feeder behave over time, the voltage history of a node located midway on Barranca Mesa, the branch below location B3

in Fig. 2, is shown in Fig. 24b. This location gives an interesting insight into the physical behavior of the system due to its distance from the source node and the fact that it is the heaviest populated area in the feeder, consuming more than a third of the total load. Its profile reveals that the voltage at this location remained between ANSI C84.1 voltage specifications and experienced a maximum voltage of 103.7% and a minimum of 99.5% of the nominal voltage.

B. Residential DR

In this case, residential demand response is used to control a TCL system of 1600 AC units with a maximum power capacity of 400kW. The system is set to respond to a control signal generated from solar irradiance and operates between 9:00AM-5:00PM, the period of the day when the sun is the most intense. Note that the loads associated with the critical infrastructure and the power provided by the DERs remains as stated in the baseline case including the battery.

The demand response performance is illustrated in Fig. 25 with the aggregated power consumption of the AC units represented in purple, while the filtered AC consumption and the filtered solar irradiance are represented in yellow and green, respectively. Here, the filtered solar irradiance is the desired output of the system, and thus, is used as the demand response control signal. Similarly, the filtered AC consumption represents the system's response to the signal. The filtering of these components is done using the bandpass recursive time series digital filter discussed in Section 2.2.2. It allows for a proper comparison between the control signal and the output of the system within the same frequency domain. It is evident from inspecting Fig. 25, that the response of the system is in strong alignment with the desired output, indicating that the demand response control was successful during the control period of this simulation. However, a closer inspection of the filtered AC consumption at 12:00, shows an example of a case when the TCL system runs out of

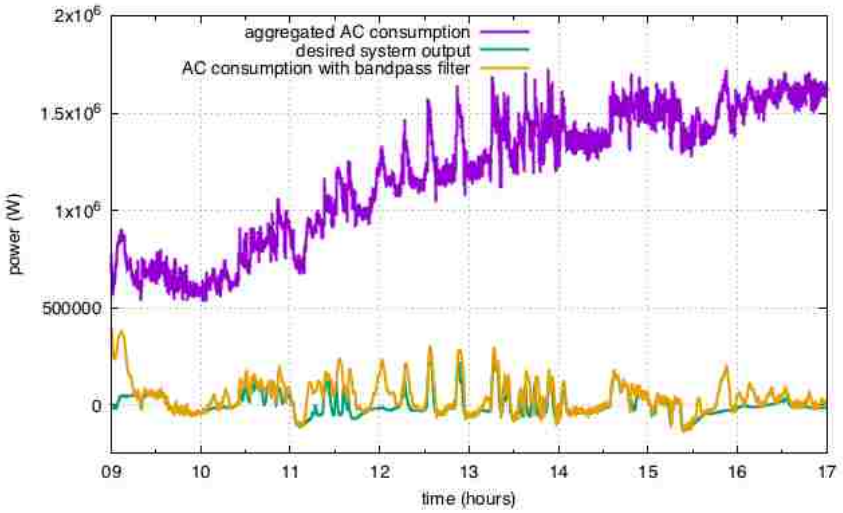

Fig. 25: Aggregated DR successfully following signal from solar irradiance between 9AM-5PM. Total consumption of 1600 AC units in purple, filtered load and signal in yellow and green, respectively.

energy and is unable to meet the control demand. This happens because there are not enough AC units available to turn OFF, in other words, the AC units that are ON cannot change states. Doing so, would cause them to violate their thermostat setpoint compromising the quality of service to the customer.

The effects of demand response on the distribution feeder and on the overall residential power consumption are observable in Fig. 26 compared to the baseline results. As expected, both figures show identical profiles from midnight until 09:00 when demand response is enabled. Once DR is active, the load reduction on the transformer is clear, especially notable around 09:20 when demand response reduces the load by almost 1MW. The strong agreement in behavior from figures 25 and 26, indicates that the substantial peak at this time of the day was a direct response of AC units reacting to the rising outside temperature. The absence of this peak during demand response further supports this observation since the AC power consumption is the only varying element between both scenarios, indicating that this peak was successfully eliminated by demand response. Furthermore, it is also evident from Fig. 26

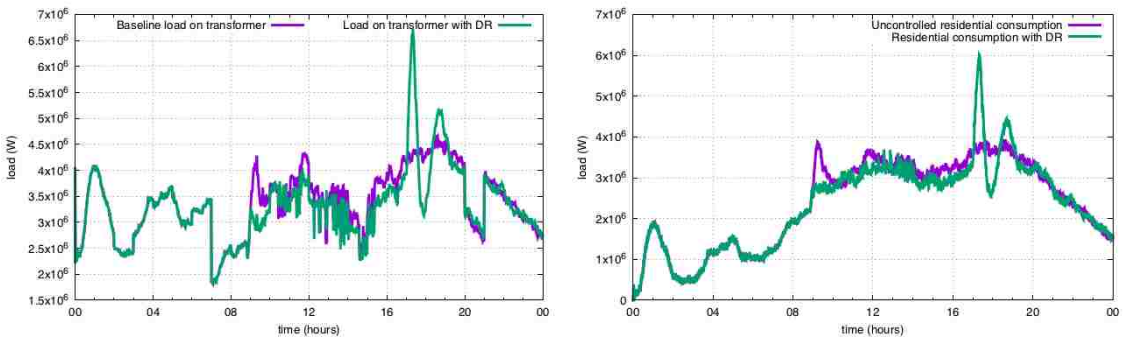

(a) Substation Transformer load comparison between baseline and DR case.

(b) Residential load comparison between baseline and DR case.

Fig. 26: Circuit 16 load comparisons between baseline and DR cases. (a) Load on the substation transformer under uncontrolled conditions & with demand response. (b) Residential load resulting from baseline and demand response cases.

that the load was reduced accordingly in response to the solar irradiance signal for the duration of the DR control period. However, an interesting and unexpected phenomenon is observed after demand response is disabled at 17:00. The residential load experiences high-amplitude oscillations, known as the demand response rebound effect, for about two hours before settling to a normal load level with minor oscillations. The first oscillation immediately after stopping DR, results in a spike of roughly 2.5MW in a matter of 20 minutes, causing a fuse in the middle branch of the feeder to blow. In addition, a substantial voltage drop across the feeder was experienced in response to the massive load increase. The voltage profiles shown in Fig. 27, show a comparison between the voltage at each node across the feeder as a function of distance from the source node for the baseline and the case with demand response at the time of peak load (17:20). It is evident from comparing Figs. 27a and 27b that this phenomenon also had substantial effects on the voltage levels in the entire Circuit 16 feeder. In fact, phase C experienced a voltage of 95.32% of nominal; within 0.32% of violating the ANSI C84.1 voltage range specifications denoted by the red lines in Fig. 27.

This phenomenon results from the deactivation of the demand response control. At

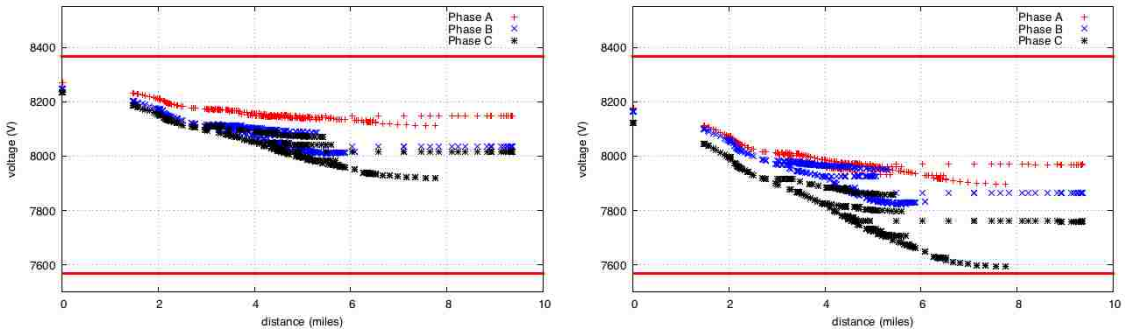

(a) Circuit 16 voltage profile in baseline case. (b) Circuit 16 voltage profile in DR case.

Fig. 27: Circuit 16 voltage profiles as a function of distance from the source node recorded at 17:20 from baseline and DR cases. Note the red lines denoting the voltage range limits from ANSI C84.1 specifications. (a) The voltage profile during typical operating conditions exhibits normal operating voltage levels. (b) Voltage profile during a wild load oscillation with a peak of about 2.5MW experienced 20 minutes after DR was disabled.

the end of the control period, almost all of the AC units available for load reduction in the system are used, as indicated by the green curve on Fig. 28. In other words, most of the AC units in the system are either already OFF, or they are not allowed to turn OFF due to their temperature setpoints and constraints. Therefore, once the AC units are not required to adhere to a control signal, they return to a relaxed state within a comfortable temperature between their setpoints, similar to when the load is removed from a loaded spring. Fig. 28 shows that while the number of units available to turn OFF remains low and relatively constant throughout the control period, the number of units ON increased. It also shows that right before DR was disabled, a total of about 400 AC units were active and 1200 were OFF, 800 of which were available to turn ON. However, only a small fraction of the active units were available for load reduction. After DR was disabled, 644 of the 800 units available to turn ON became active. At the 4kW consumption per unit defined in the model, this amounts to a total load of 2.576MW, which is in strong agreement with the spike observed in Fig. 26. Moreover, this occurrence coincided with the maximum

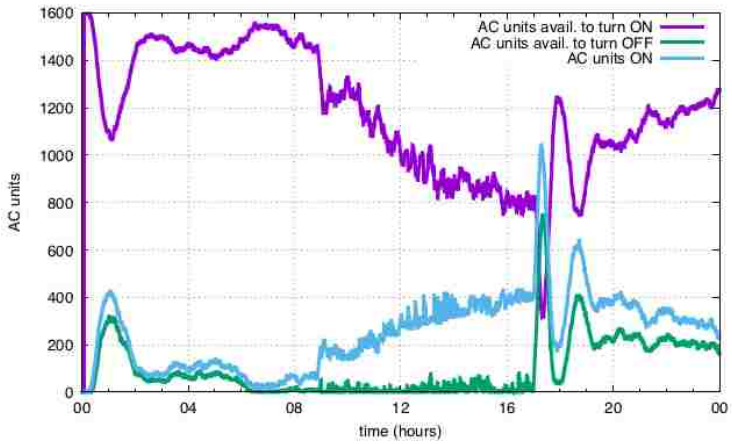

Fig. 28: States of the AC units participating in demand response. The graph shows the number of AC units that are available to turn ON and OFF, as well as the number of AC units that are ON throughout the day. Note the spike of units ON after DR is disabled after 17:00.

outside temperature, which may have led to an amplified effect of this behavior. The fact that a high number of units became active at the same time means that most of them, in addition to the ones that were already ON, reached a comfortable temperature and became inactive at a similar time, thus, leading to the substantial drop in load observed at 18:00. However, the distribution of setpoint values among thermostats caused their synchronized operation to normalize to a relatively constant number of operating units, similar to the behavior experienced during the first three hours the simulation. A potential solution to this behavior is to ramp-down the removal of demand response control, or to disable control by grouping AC units and releasing them in waves. Another option is to implement load forecasting and more sophisticated controls such as an MPC, in order to minimize or eliminate the negative effects of demand response.

The effects imposed on the voltage experienced by Barranca Mesa after demand response was disabled are shown in Fig. 29 in comparison to the voltage history of the same point obtained during the baseline case. At the time of peak consumption after

demand response was disabled, Barranca Mesa experienced a voltage of 96.49% of nominal, a difference of 3% from the case under normal operating conditions. The results from this simulation further highlight the value of combining control scheme with power flow properties into a single co-simulation. For instance, simulating demand response without a power flow simulator would neglect the effects of rebound on the electrical infrastructure. It would be impossible to examine potential voltage violations or understand the locations of faults in the system. This tool grants the ability to test the performance of new control algorithms while considering aspects of human behavior, and to analyze their effects on the grid. It also provides the ability to pinpoint locations where faults and violations might occur, allowing to devise potential solutions based on the insight gained from the results, and to test those solutions using the same simulation framework.

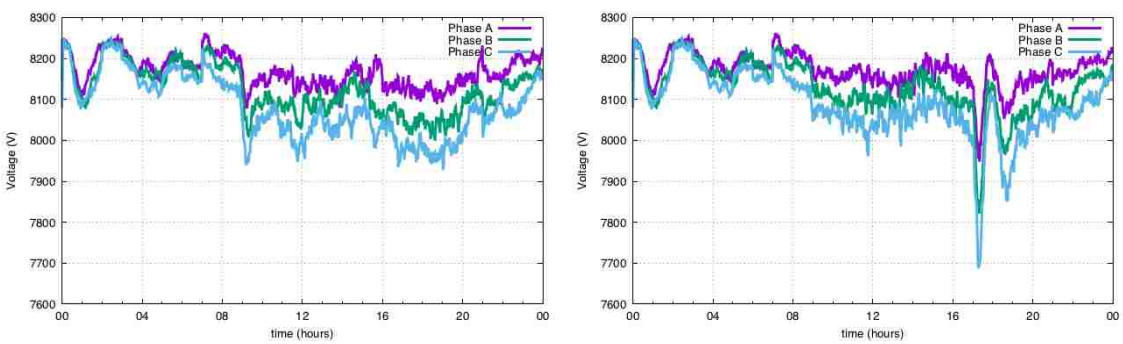

(a) Voltage history of a branch in baseline case.

(b) Voltage history of a branch in demand response case.

Fig. 29: Comparison of voltage history of Barranca Mesa branch during baseline and DR cases. (a) Voltage history of a node located midway on the branch below location B3 under baseline conditions. (b) Voltage history of the same node on the branch under location B3. Note the substantial voltage drop at 17:20 after DR is disabled.

C. Battery control & DR

Demand response, as well as the battery controller were used together in this case to maintain the load at the substation transformer at or below 3.5MW. As in Section 4.1.1.B, the initial battery SOC is assumed to be 0.9 at midnight from control actions performed the previous day. The controller dispatches PV and the battery if necessary to keep the load on the transformer at the target value. The same control of 1600 AC residential units via demand response is used here with the same control schedule between 9:00AM-5:00PM, with the same control signal created from solar irradiance. As a consequence, the residential load remained identical to the load seen in Section 4.1.2.B including the oscillating phenomenon experienced after demand response is disabled. Here, the main grid supplies power to the feeder, while the demand response lowers the residential consumption according to its control signal,

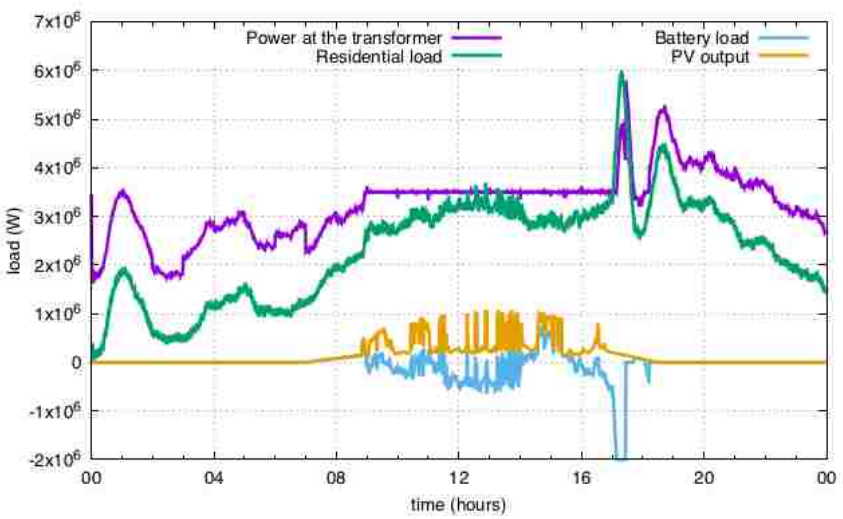

Fig. 30: History of load on substation transformer during support by the battery controller and residential demand response. Residential load, battery dispatch, and PV generation are also shown for comparison. Battery control successfully dispatches PV and battery power to maintain the load on the feeder at the target load of 3.5MW, until the load increase experienced after the DR control period drains the battery to its reserve SOC of 25%.

and the controller minimizes the difference between the load on the transformer and the target load by dispatching PV and battery resources.

The overall load on the substation transformer in comparison to the residential and battery load profiles including PV power generation are shown in Fig. 30. Note that the PV output and the residential profiles are identical to cases in the previous sections. However, battery output is varied by the controller according to the power needed on the feeder to maintain the load at the target level. Fig. 30 shows that the load on the transformer remains below the predefined threshold of 3.5MW until 09:00 when the load begins to increase as a result of the typical morning peak. From this point forward, the controller effectively dispatches PV and battery resources in order to maintain the load at the target value for a period of over eight hours. However, the increased power consumption associated with the synchronized activation of the 644 AC units after demand response is disabled overpowers the support DER support. Dictated by the lack of sunlight at this time of the day, only the battery is dispatched, in order to counter the effects of the sudden load increase experienced by the transformer. The battery sends power to the feeder at 2MW for almost 20 minutes until it reaches its reserve SOC of 25%, allowing the load to only reach about 5MW. This is more than 1.5MW less than in the previous case at the same point in time. Once the battery reaches its reserve SOC level and becomes inactive, the load on the transformer spikes up again due to the high consumption of the AC units that remains. However, the highest load experienced during the entire event is 1MW less than that experienced during the case without battery control.

While the DERs in this scenario were not enough to mitigate the effects of the demand response rebound, the action of the battery was enough to lessen the disruption to the feeder voltage levels at the time of highest consumption that took place at 17:20. The voltage at every node in the residential demand response only case discussed in the previous section, in comparison to the voltage profile obtained in this scenario are shown in Fig. 31. With battery support managed by the controller, the

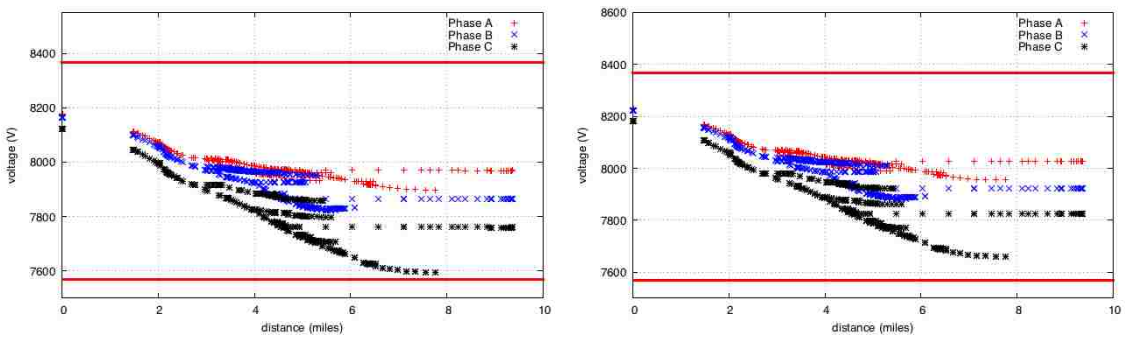

(a) Circuit 16 voltage profile in demand response case.

(b) Circuit 16 voltage profile in DR & battery control case.

Fig. 31: Circuit 16 voltage profiles as a function of distance from the source node recorded at 17:20 from DR-only and DR & battery control cases. Note the red lines denoting the voltage range limits from ANSI C84.1 specifications. (a) Voltage profile during consumption spike of about 2.5MW experienced 20 minutes after DR was disabled during DR only case. (b) Voltage profile at same time of peak consumption after DR was disabled with battery support. The lowest voltage experienced with battery support was 96.14% compared to 95.32% during DR only case.

lowest voltage value experienced on the feeder was recorded on phase C at 96.14% of the nominal voltage. This is a 0.82% difference from the lowest voltage value observed in Section 4.1.2.B (95.32% of nominal). A similar impact was felt in the Barranca Mesa branch during demand response rebound with battery support, as shown in Fig. 32. While the voltage drop remains substantial, it is noticeably less drastic when compared to the unsupported case. Furthermore, the lowest voltage value experienced at this location on the feeder was recorded at 97.33% of the nominal voltage, a difference of 0.84% and 2.17% from the case with demand response only and the case under normal operating conditions, respectively.

The results of this simulation clearly indicate that it is possible to reduce the disruption caused by demand response rebound, or other similar events, by employing battery support with more sophisticated controls. For example, a load forecaster or MPC would take data from this simulation such as the load on the feeder and battery SOC, to predict similar behavior in the future and optimize the use of resources to

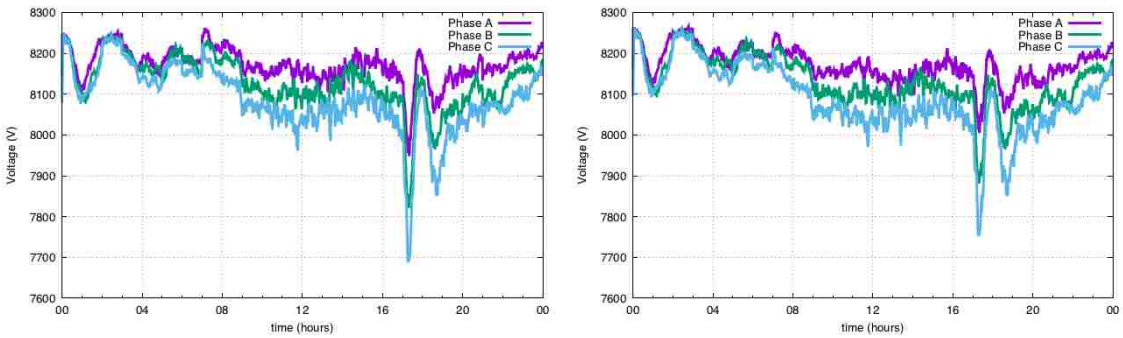

(a) Barranca Mesa voltage history during DR-only case.

(b) Barranca Mesa voltage history during DR & battery control support.

Fig. 32: Comparison of impact imposed on Barranca Mesa voltage after deactivation of demand response during DR-only and DR & battery control cases. (a) Barranca Mesa voltage with only demand response support. Voltage levels drop to 95.32%, almost violating ANSI C84.1 voltage limit specifications. (b) Voltage experienced by Barranca Mesa during support by demand response and battery control. The voltage drop in this case is remarkably less drastic due to added support by the battery.

minimize the effects associated with disruptive events. Additionally, tools such as Fig. 31 provide the ability to identify sections of a feeder at risk of violating voltage limit specifications. In this case, it is evident that the location of the feeder corresponding to the nodes on phase C that experience the lowest voltage levels, would benefit from injection of reactive power or additional support from DERs. In fact, it is theoretically possible to generate a constant voltage profile across the feeder with enough control and DER support at strategic locations on the system. This valuable insight would remain obscure without the ability of this framework to analyze how power flow is affected by certain control schemes.

4.2 Real-time simulation

The performance of the real-time simulation setup described in Section 3.2 is introduced here. The residential load generator is used to simulate the residential load of the 200-home community located in Mesa del Sol. Real-time residential demand response capabilities are enabled by the load aggregator, who controls the aggregated power consumption of the air conditioner system in each home using control signals provided by a microgrid controller. The controller optimizes the use of all resources in the system discussed in Section 1.3 for the minimum cost of operation. However, the requests from the microgrid controller to the aggregator, were not very eventful. The setpoint for the TCL generally remained at zero due to the existence of more cost-effective resources that were used to satisfy the demand. Therefore, a synthetic signal was created to examine the real-time demand response behavior of the system.

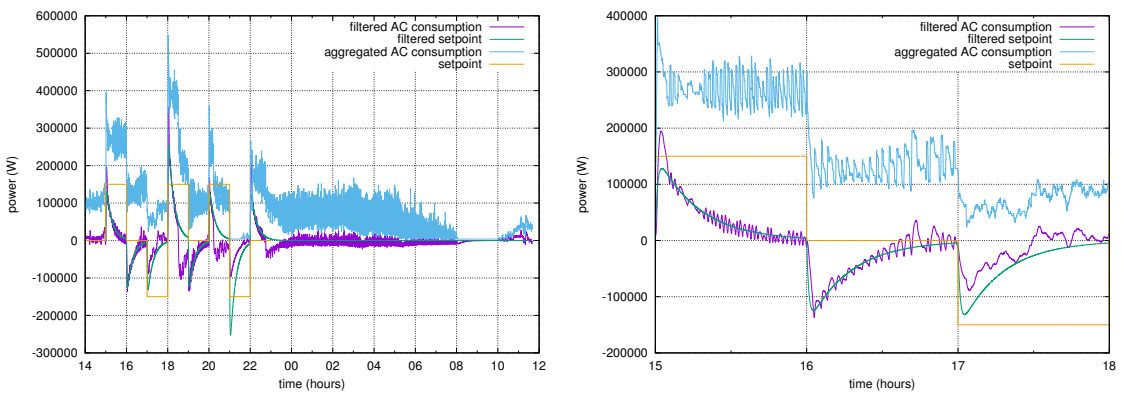

(a) Performance of DR control during a 22-hour test period.

(b) Closeup at DR performance during a 3-hour period.

Fig. 33: Performance of real-time residential demand response during a test period of 22 hours using a synthetic control tracking signal. (a) Aggregated TCL reaction to demand response control. The response of the system is as expected; the filtered TCL response tracks the filtered signal over the course of the entire control period. (b) Focused view at the performance of DR control during 3 hours of operation. Note the inability to track the control signal accurately after hour 17, due to lack of available resources.

The signal created varies hourly with requests to decrease and increase power consumption at different magnitudes for a period of nine hours. The setpoint remains at zero for the following 13 hours, for a total test period of 22 hours as shown in Fig. 33a. It is evident from inspecting Fig. 33a, that the TCL resource effectively follows the filtered control signal for the majority of the test period. However, in certain cases the TCL's "thermal battery" runs out of energy preventing the TCL from fulfilling the control request, similar to 12:00 in Fig. 25. This behavior is observed in hour 17:00 of Fig. 33b where a requests to decrease power consumption cannot be satisfied as it would force house temperatures to increase above the deadband compromising the customer's comfort. This drawback is common among demand response schemes in which residential customers participate with the condition that

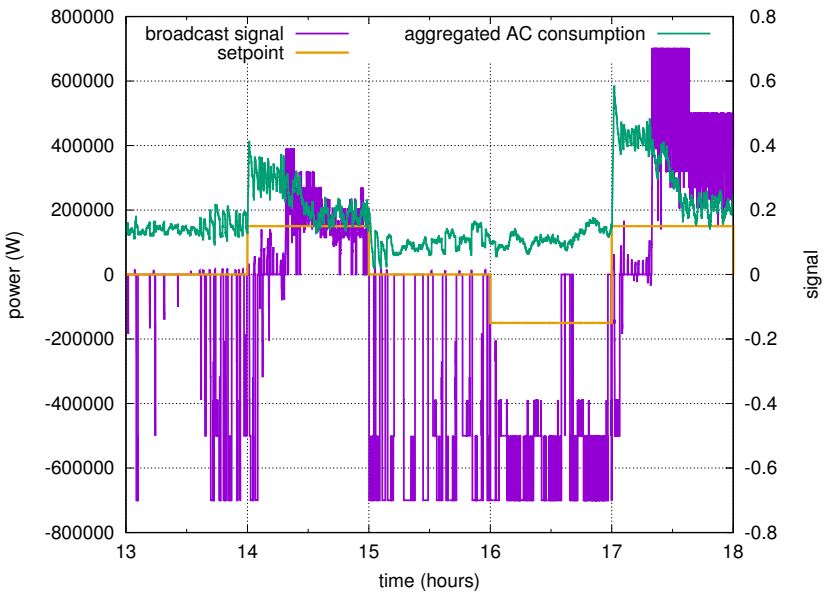

Fig. 34: Example of control signals and response of residential demand response control. The graph shows the synthetic control signal sent by the microgrid controller to the aggregator, with the associated control signal broadcasted by the aggregator to the participating TCL units. The aggregated response of the system to the broadcast is represented by the green line. Note the shifts of consumption in response to the control signals.

quality of service be the priority. Moreover, the negative impact of this limitation could be minimized with a more complex policy where financial incentives are offered to flexible customers who are willing to accept remote controllability of their temperature setpoint.

The relationship between the control signal issued by the microgrid controller, the signal broadcasted by the load aggregator to the TCL units, and their response is illustrated in Fig. 34. It shows the setpoint control signal issued by the microgrid controller, the broadcast signal sent by the aggregator to the TCLs and their aggregated response of the system. A closer inspection of Fig. 34 reveals the relationship between the data exchanged by the controller and the aggregator, and the response of the TCLs. The broadcast signal displays a heavy bias towards the control signal as expected, especially noticeable during hours 14, 16, and 17.

The real-time results from the simulation, as explained in Section 3.2, are sent to a central database by each model. These results are showcased at a publicly available

Fig. 35: Simulated residential load used in the MRI real-time simulation. The load is simulated by the load generator and used as a proxy to represent the power consumption of a 200-home residential community in a microgrid controller study performed on a feeder in the Mesa del Sol community in Albuquerque, NM. This figure was generated using real-time data obtained from the study's database and posted on a UNM website located at 129.24.68.15.

research website located at the IP address 129.24.68.15. Results include the total load of the virtual Mesa del Sol residential community and other resources on the system, such as PV arrays, batteries, fuel cell, and other components that lie beyond the scope of this work. A snapshot of the community's power consumption obtained from the website is shown in Fig. 35. The graph illustrates the load history of the residential community in a period of two hours between 7:30-9:30AM with five-minute resolution. Note that the graph shows the morning ramp-down that takes place after people leave their homes for work.

Here, the ability to simulate real-life communication over a network between residential infrastructure and a control center was successfully demonstrated. It was shown that the physical separation of these components is negligible and that remote control of residential loads is possible using the internet and internet-enabled appliances such as smart thermostats. The successful integration of these components into the real-time microgrid controller simulation with HIL equipment demonstrates the ability to use this framework as a powerful tool to conduct research that considers the effects of control algorithms on existing electric infrastructure.

Chapter 5

Discussion & Conclusions

5.1 GridLAB-D co-simulation

The main purpose of the GridLAB-D co-simulation framework described in this thesis, is to provide the means for collaboration between a diverse group of models that are originally standalone, in order to achieve a combined realistic representation of a power distribution system's behavior. This is accomplished by providing a platform with a direct line of communication and mutual synchronized collaboration between models. The platform enables dynamic results that evolve in response to the interactions between co-simulators, while each model operates within its individual domain. Here, the co-simulation framework combines the GridLAB-D power flow engine to simulate the distribution of power within a feeder, with the abilities of a residential load generator to accurately simulate the residential load of a community based on human behavior models, and to perform demand response control of household appliances using a load aggregator.

The Matlab interface co-simulation setup demonstrated the potential of the framework by successfully incorporating human behavior models with the GridLAB-D

power distribution simulator. It set a benchmark performance of the models, and also provided proof of concept for the GridLAB-D built-in battery controller and demand response capabilities of the load aggregator. However, the Matlab interface was unreliable and it permitted only a small amount of data to be transferred between models. It also required a vast amount of computational resources that resulted in extremely slow execution times. These limitations made clear the necessity for a more robust and powerful replacement of the Matlab interface.

Developed by Pacific Northwest National Laboratory, FNCS was designed with the sole purpose to serve as a co-simulation interface for software that cannot operate satisfactorily together under typical computer conditions. FNCS enabled the ability to transfer more than 200 times the amount of data between co-simulators, twice as fast execution times, and more efficient and reliable simulation executions. The results from the FNCS baseline simulation showed similar results to the uncontrolled scenario obtained with the Matlab interface, indicating that the simulated infrastructure behaved within physical limits of the system. Furthermore, the enhanced capabilities of the FNCS interface demonstrated the ability to employ demand response control of an entire feeder in combination with utility-scale battery controls with the objective to reduce the stress on the power distribution infrastructure. While the design of a demand response infrastructure is beyond the scope of this work, the capabilities to simulate control of an aggregated system of household appliances was successfully demonstrated. Also, although the unexpected behavior associated with the deactivation of demand response control was analyzed and discussed, an in-depth examination of this phenomenon is aligned with the development of demand response infrastructure, and thus, it too lies outside of the scope of this thesis. However, the simulation framework showcased here has the potential to be a powerful tool in the development and design of future demand response algorithms and infrastructure, in an effort to improve the resiliency of the power grid.

5.2 Real-time simulation

The development of modern microgrid controllers, such as the one demonstrated in the Mesa del Sol area in Albuquerque, NM, requires the ability to test the system's behavior under different circumstances. Testing is especially crucial in the development of demand response algorithms due to the risks involved in deploying systems that have the potential to disrupt the quality of service experienced by the customers if not designed properly. This requires testing communities with real smart appliances and the infrastructure necessary to conduct test. The high upfront cost and risks associated with this type of testing method means that it is often avoided. In addition, since this technology is relatively young, in most cases the communities are not equipped with demand response capabilities. A more cost effective and convenient method is to perform computer simulations using virtual models to analyze the behavior of a system before physically deploying resources, as was done for the Mesa del Sol microgrid controller discussed in Section 1.3.

Here, the residential load generator and aggregator were incorporated into a real-time simulation that tests the capabilities of the micogrid controller to manage and optimize the use of DERs on a distribution feeder. The simulation incorporates both virtual and real elements including a commercial-scale microgrid, utility-scale PV and BESS, residential demand response, and VPP and MPC for optimization. The load generator was used as a proxy for the 200-home Mesa del Sol community with simulated demand response capabilities implemented with the use of the aggregator. Both of these models were successfully installed onto individual Raspberry Pi computers to simulate the physical separation of the residential infrastructure and load aggregator in real life. In addition, communication between the models and the microgrid controller database was employed via the Internet, successfully replicating communication of similar devices in the field. The results discussed in Section 4.2 show the effective performance of the aggregated control and response of the TCL

system to the control signal issued by the microgrid controller for a period of nine hours. The HIL real-time simulation platform showcased in this thesis demonstrates the ability to incorporate existing distributed energy resources within a partially virtual environment that fosters the development of future infrastructure essential for modern microgrids. Most importantly, this platform enables the ability to test new microgrid controller algorithms without risking equipment damage or quality of service to the customer. At the same time, it also provides the ability to gain valuable insight related to the introduction of new elements and their impact on existing electric infrastructure. These capabilities are essential in the development of microgrid technology that will improve grid resilience by providing critical services in the event of severe disruptive events.

5.3 Future Work

While the models and simulation frameworks described in this work generated successful results, improvements to some of the components is still necessary to make them more effective. These include, enhancements to the residential load generator that simulate changes in human behavior due to outside influence. In this case, behavior demand response will be introduced to replicate economic incentives granted by the utility to customers participating in load reduction programs. Additional control testing is also required alongside the residential load aggregator to determine an appropriate solution to rebound behavior experienced after demand response control was disabled. A potential solution to this problem is the introduction of a load forecaster or MPC to predict disruptive events while optimizing the use of resources. The addition of a communication infrastructure simulator to the FNCS co-simulation framework is also on the scope of future work for this research project. The ability to combine aspects of communication to the framework will allow to consider char-

acteristics of demand response communication between smart appliances and the aggregator that are currently neglected. This addition will also allow to consider aspects of cybersecurity associated with the power grid.

Currently, the GridLAB-D built-in battery controller is only capable to charge from power generated by an accompanying PV array through a shared inverter. Therefore, in order to represent real-life system configurations and capabilities, the battery controller should include the ability to charge from the grid as well as directly from PV through individual inverters. However, its performance illustrated in Sections 4.1.1.B and 4.1.2.C confirmed that the control logic used for dispatching PV and battery resources is effective.

The real-time communication method between the load generator and the aggregator via the Internet also requires improvements. Considered "brute force," the current communication method is secure, but very inefficient. The models exchange information by copying and pasting files between machines over a network, and are opened and read on every time-step. This could be improved by implementing a standardized communication protocol such as DNP3, which also provide high levels of security, comparable to the Secure Copy command line tool.

5.4 Conclusions

If he were alive, Thomas Edison would recognize most of the power grid infrastructure used today despite the technological advances made to electric power infrastructure over the last 150 years. For instance, most of the world still distributes electricity from power stations to its customers over power lines mounted on wooden electric poles, much like in the days of Edison. However, the power grid is currently undergoing a major transformation due to recent technological advances, especially in the areas of computer science, telecommunications, and electrical engineering, as well as

an increase penetration of distributed energy resources that is driving the greatest leap forward in the history of the power industry. This is forcing industry stakeholders to adapt new business models that modernize the power grid infrastructure in a way that considers endusers not only as customers, but also as potential contributors to the operations and management of power grid resources. The increasing penetration of residential DERs, as well as emerging internet enabled appliances such as smart thermostats, indicates that residential communities can be considered low-scale power plants with capabilities to offset loads using demand response on their appliances. However, the standards and guidelines in existence today require further research and development before full deployment of such infrastructure is feasible. Also, the development of this technology demands studies that consider human behavior to examine the effects of the human factor in distribution-level load management and demand response.

The work presented in this thesis highlights two co-simulation platforms that combine distribution system simulators with a residential load synthesis framework based on human behavior models. The results obtained from the various simulations performed with each platform successfully established their performance and showed their potential as tools to conduct research in the efforts to develop a modernized power grid infrastructure. Even though improvements to the models are necessary for enhanced behavior of the systems, the co-simulation platforms' performance was effective, unlocking endless possibilities for future co-simulations that combine additional models regardless of programming language or compatibility. As the electric power industry evolves with innovation, it will require similar innovative methods to study and develop the technology necessary to create the power grid of the future, one that is more efficient, reliable, resilient, and smart.

Appendices

Appendix A

Matlab interface models & code

The models and related files used in the co-simulations performed with the Matlab interface can be found in the following github repository.
https://github.com/victorayon/masters-thesis/tree/master/Matlab-interface

Appendix B

FNCS interface models & code

The models and related files used in the co-simulations performed with the FNCS interface can be found in the following github repository.
https://github.com/victorayon/masters-thesis/tree/master/FNCS-interface

Appendix C

Real-time simulation models & code

The models and related files used in the real-time simulation can be found in the following github repository.
https://github.com/victorayon/masters-thesis/tree/master/Real-time-simulation

References

[1] V. Ayon, M. Robinson, A. Mammoli, A. Fisher, and J. Fuller, *"Integration of Bottom-up Statistical Models of Loads on a Residential Feeder with the GridLAB-D Distribution System Simulator, and Applications,"* Innovative Smart Grid Technologies (ISGT) Europe, IEEE, Sep. 2017.

[2] V. Ayon, M. Robinson, and A. Mammoli, *"Simulation of Real-time Demand-response to Support a Residential Distribution Feeder in Microgrid Mode,"* Grid of the Future, CIGRE US National Committee, Oct. 2017.

[3] D. P. Chassin, J. C. Fuller, and N. Djilali, *"Gridlab-d: An agent-based simulation framework for smart grids,"* Journal of Applied Mathematics, vol. 2014, 2014.

[4] *"Matlab link,"* `http://gridlab-d.shoutwiki.com/wiki/Matlab_link`, accessed: 2017-09-26.

[5] *"Chapter 9 - Interfacing with External Software,"* `http://gridlab-d.shoutwiki.com/wiki/Chapter_9_-_Interfacing_with_External_Software`, accessed: 2017-08-18

[6] D. Ton and J. Reilly, *"Microgrid Controller Initiatives,"* IEEE Power & Energy, no. 4, vol. 15, pp. 24-25, Jun. 2017

[7] A. Mammoli *et al.*, *"Validation of Measures to Expand Applications of Micro-Grid Control System,"* University of New Mexico & Sandia National Laboratories, for Mitsubishi Research Institute, Sep. 2017.

[8] *"Milsoft engineering analysis WindMil,"* `https://www.milsoft.com/utilitysolutions/upgrades/milsoft-engineering-analysis-ea-windmil`, accessed: 2017-03-11.

[9] *"C84. 1-1982,"* American National Standard for Electric Power Systems and Equipment-Voltage Ratings (60Hz), aNSI Standard.

[10] A. Mammoli, M. Robinson, V. Ayon, M. Hombrados-Herrera, M. Martinez-Ramon, *"A simulation framework to develop control and forecasting tools for aggregated residential energy resources,"* Applied Energy, 2017, under review.

[11] R. Hendron and C. Engebrecht, *Building America house simulation protocols.* National Renewable Energy Laboratory Golden, CO, 2010.

[12] D. Fischer, A. Härtl, and B. Wille-Haussmann *"Model for electric load profiles with high time resolution for german households,"* Energy and Buildings, vol. 92, pp. 170-179, 2015.

[13] J. Widén, A. M. Nilsson, and E. Wäckelgård, *"A combined markovchain and bottom-up approach to modelling of domestic lighting demand,"* Energy and Buildings, vol. 41, no. 10, pp. 1001-1012, 2009.

[14] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *"Digital Filtering in the Time Domain,"* in Numerical Recipes in Fortran, 2nd ed. Cambridge, England: CUP, 1986-1992, ch.13, sec. 5, pp. 551-556

[15] *"FNCS-Tutorial/demo-large/,"* `https://github.com/GridOPTICS/FNCS-Tutorial/tree/master/demo-large`, accessed: 2017-08-20.

[16] *"1. WHAT IS A RASPBERRY PI?"* `https://www.raspberrypi.org/help/faqs/#introWhatIs`, accessed: 2017-08-05.

[17] *"12 scp command examples to transfer files on Linux"* `https://www.linux.com/blog/12-scp-command-examples-transfer-files-linux`, accessed: 2017-08-05.

[18] *"SSH COMMAND"* `https://www.ssh.com/ssh/command/`, accessed: 2017-08-05.

[19] O. Lavrova, F. Cheng, S. Abdollahy, H. Barsun, A. Mammoli, D. Dreisigmayer, S. Willard, B. Arellano, and C. Van Zeyl, *"Analysis of battery storage utilization for load shifting and peak smoothing on a distribution feeder in new mexico,"* in Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES. IEEE, 2012, pp. 1?6.

[20] F. Cheng, S. Willard, J. Hawkins, B. Arellano, O. Lavrova, and A. Mammoli, *"Applying battery energy storage to enhance the benefits of photovoltaics,"* in Energytech, 2012 IEEE. IEEE, 2012, pp. 1?5.