Pacific Northwest
NATIONAL LABORATORY

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

Coordination and Control of Flexible Building Loads for Renewable Integration

Demonstrations using VOLTTRON™

October 2016

H Hao   S Huang
G Liu   S Katipamula

Coordination and Control of Flexible Building Loads for Renewable Integration

Demonstrations using VOLTTRON™

H Hao   S Huang
G Liu   S Katipamula

October 2016

Pacific Northwest National Laboratory
Richland, Washington 99352

Abstract

Renewable energy resources such as wind and solar power have a high degree of uncertainty. Large-scale integration of these variable generation sources into the grid is a big challenge for power system operators. Buildings, in which we live and work, consume about 75% of the total electricity in the United States. They also have a large capacity of power flexibility due to their massive thermal capacitance. Therefore, they present a great opportunity to help the grid to manage power balance. In this report, we study coordination and control of flexible building loads for renewable integration. We first present the motivation and background, and conduct a literature review on building-to-grid integration. We also compile a catalog of flexible building loads that have great potential for renewable integration, and discuss their characteristics. We next collect solar generation data from a photovoltaic panel on Pacific Northwest National Laboratory campus, and conduct data analysis to study their characteristics. We find that solar generation output has a strong uncertainty, and the uncertainty occurs at almost all time scales. Additional data from other sources are also used to verify our study. We propose two transactive coordination strategies to manage flexible building loads for renewable integration. We prove the theories that support the two transactive coordination strategies and discuss their pros and cons. In this report, we select three types of flexible building loads—air-handling unit, rooftop unit, and a population of water heaters (WHs)—for which we demonstrate control of the flexible load to track a dispatch signal (e.g., renewable generation fluctuation) using experiment, simulation, or hardware-in-the-loop study. More specifically, we present the system description, model identification, controller design, test bed setup, and experiment results for each demonstration. We show that coordination and control of flexible loads has a great potential to integrate variable generation sources. The flexible loads can successfully track a power dispatch signal from the coordinator, while having little impact on the quality of service to the end-users.

Summary

Decarbonizing the global energy system is one of the most important endeavors of our time. It is widely accepted that a sustainable energy future will increasingly rely on renewable energy sources such as wind and solar. However, these variable generations are volatile, intermittent, and uncontrollable. The vast integration of variable generation sources brings a significant amount of uncertainty to the power grid, and presents a daunting challenge to current power system operation and control. Maintaining the power balance with uncertain renewables is one of the most important problems in smart grid research. Traditionally, the uncertainty of renewables has been handled through supply-side (generation) reserves. However, recent studies show that deep penetration of renewables will substantially increase the procurement of additional backup generation reserves. If additional reserves are provided by fossil-fuel generators, the net carbon benefit from renewables will be diminished, generator efficiency will reduce, and an economically untenable generation mix will result. Hence, the stability and reliability of the grid will require more flexible consumption from demand-side resources.

Among various demand-side resources, buildings in which we live and work have significant impacts on energy consumption. More precisely, buildings consume about 75% of the total electricity in the United States. The massive electric power consumption and enormous thermal storage capability of buildings have a great potential for providing various grid services that are necessary to integrate uncertain renewable energies. In particular, building loads such as commercial building heating, ventilation, and air-conditioning (HVAC) systems, residential thermostatically controlled loads (TCLs), electric vehicles (EVs), and deferrable appliance loads are excellent flexible resources for absorbing the uncertainty of renewable generations. For example, by changing the supply fan speed of an HVAC system, its power deviation from baseline can be made to track the power deviation of a solar photovoltaic (PV) panel, thus mitigating its power output variability. This project aims to investigate integration of variable generation sources into the grid by aggregating, coordinating, and controlling distributed demand-side resources.

To fully exploit the potential of various classes of demand-side resources, it is essential to study optimal coordination of multiple flexible loads considering their diversified response characteristics and availability. The focus of this report is renewable integration, which aims to use various types of flexible building loads to locally absorb the uncertainty of renewables so that less power fluctuation is injected into the power grid. Based on the resources that are available, three types of flexible building loads were chosen: 1) a supply fan in an air handling unit (AHU), 2) a supply fan in a roof top unit (RTU), and 3) a collection of water heaters (WHs). Additionally, coordination and control of these three types of building load to compensate the solar generation deviation of a solar PV panel on the Pacific Northwest National Laboratory (PNNL) campus are demonstrated in this report. The overall objective of the report is to coordinate and control power consumption of flexible loads to mitigate the power fluctuation of the solar PV panel without affecting the quality of service to the end-users. Moreover, our test bed infrastructure was based on VOLTTRON™, a PNNL-developed low-cost and open-source software platform, which facilitates communicating with, and controlling flexible building loads. Although our focus is on the three specific types of building loads for the purpose of renewable integration, the developed coordination and control algorithm, and the test bed infrastructure can also be adopted to provide other building-to-grid functions.

In this report, we first analyze data from the generation output of solar PV panels on the PNNL campus. Next, for each type of the three flexible loads (AHU, RTU and WHs), system description, model

identification, controller design, test bed setup, and demonstration results are presented in this report. The results showed that coordination and control of flexible loads has a great potential to integrate variable generation sources. Tracking a power dispatch signal from the coordinator could be achieved successfully, while having little impact on the quality of service to the end-users. The main results of the technical report are summarized below:

Solar Data Analysis

The solar PV panels in this research are located at Richland, Washington. Richland has about 200 sunny days a year, which makes it a great location for installing solar PV arrays. The design power production rate of these panels is 126 kW. The VOLTTRON™ software platform was used to collect the PV panels' power production data at a time resolution of 1 minute, and store the solar data in a VOLTTRON™ historian database.
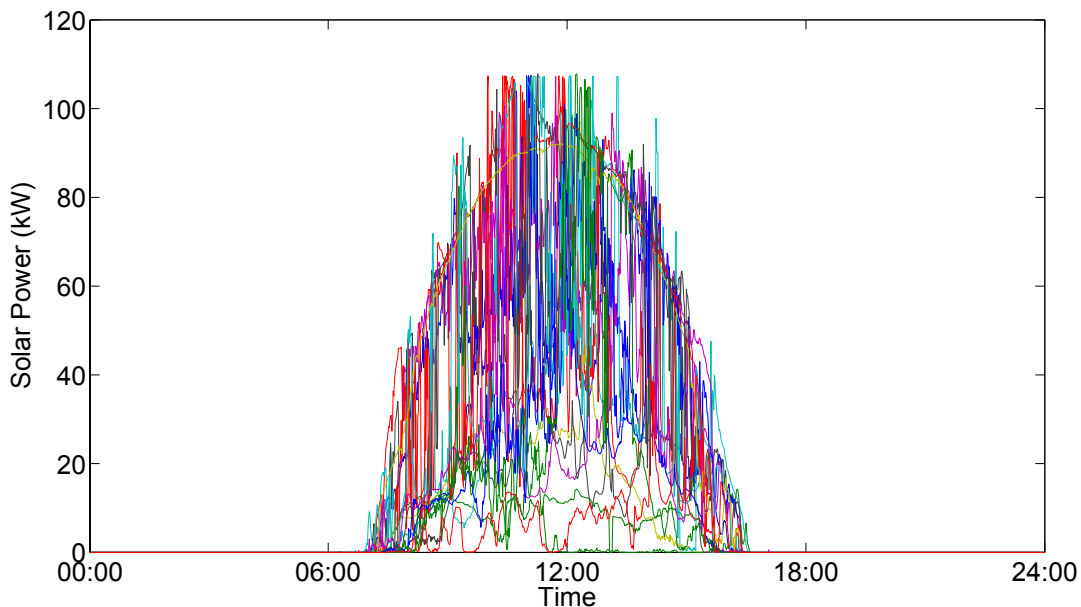

Figure S.1. Daily solar PV generation output (different colors represent different days).

The trend data showed that the solar generation output was very intermittent (please Figure S.1). On most days, its power output fluctuated frequently and abruptly due to passing clouds. The estimated power density of the solar generation data showed that the uncertainty of the solar PV output occurred at most of the frequencies, and therefore made it challenging to target a particular frequency band.

The objective of our research was to control the building loads so that their real-time power minus their baseline tracks solar generation deviation. For simplicity, we used the 15-minute trailing average of the solar generation as the solar prediction in the next time step. The results showed that the 15-minute trailing average is much smoother than the ground-truth solar power output (Figure S.2). The objective of our research is to control the building loads so that their real-time power minus their baseline tracks the solar generation deviation.
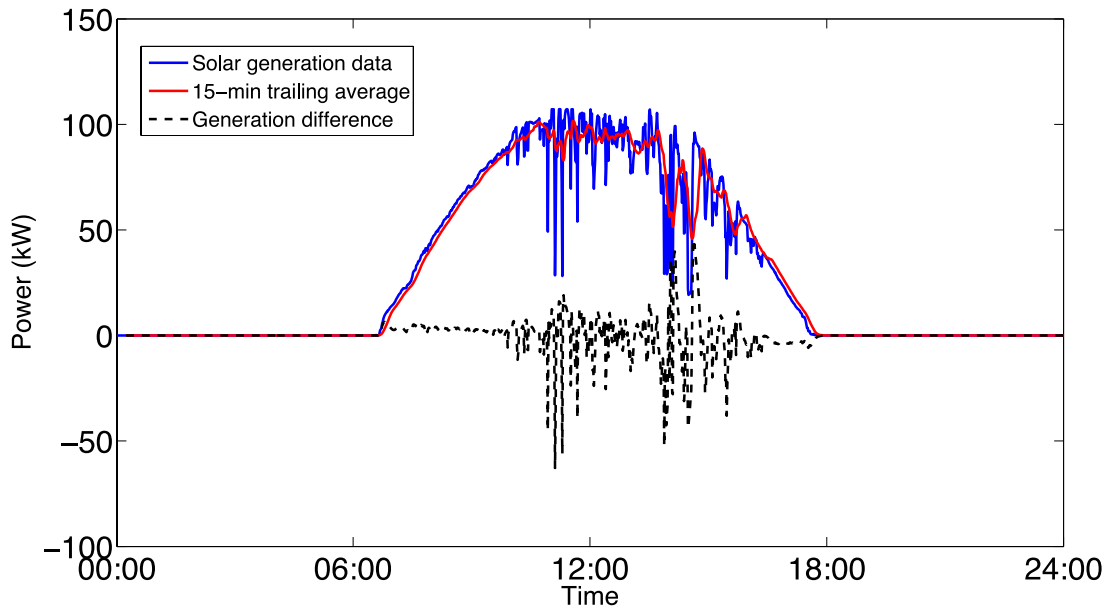
Figure S.2. Solar generation, its 15-minute trailing average, and the difference between the two for solar data on March 15, 2016.

Transactive Coordination of Flexible Building Loads

To fully exploit the flexibilities of various types of flexible loads for renewable integration, we proposed a coordination framework to manage load power consumption for renewable integration. The schematic representation of the coordination framework is depicted in Figure S.3. In this framework, there are two levels: the flexible load level and the coordinator level. The lower level includes different types of flexible loads. At each time step, loads report their forecasted baseline for the next period and power flexibility (the range of power the load is able to change from the baseline). In the upper level, a coordinator is responsible for forecasting renewable generation and implementing coordination strategies to manage the flexible loads so that their real-time power minus the baseline power can track the renewable generation deviation.
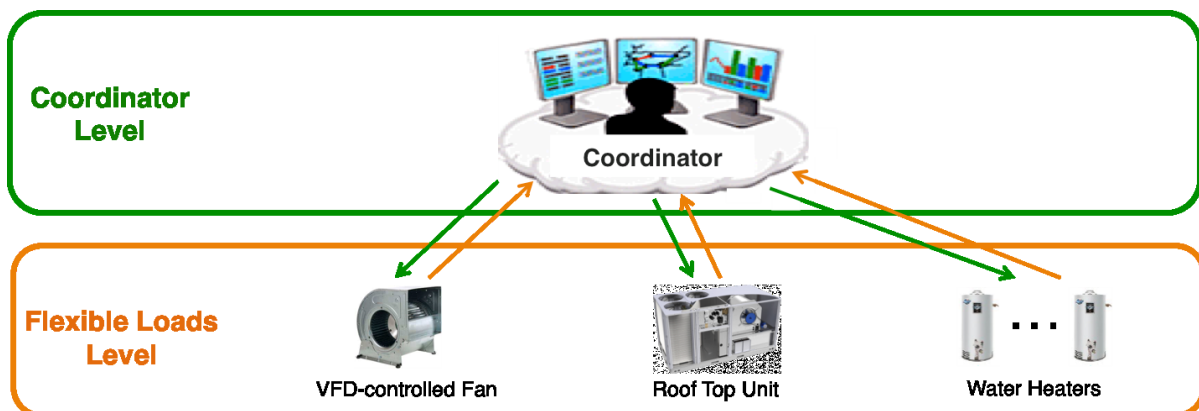

Figure S.3. Schematic of the hierarchical coordination framework.

Additionally, two distributed transactive coordination strategies (iterative method based on dual decomposition, and one-shot supply function bidding method) were proposed to engage self-interested responsive loads to provide services to the grid. We also compared the solutions obtained from the two methods with that solved using a centralized approach. We showed that these two different approaches were equivalent, and they both converged to the same optimal solution (Figure S.4). However, there was a conservation of complexity in the two transactive coordination strategies, both in terms of algorithm convergence rate and the amount of information needing to be exchanged. The iterative method based on dual decomposition required multiple iterations between the loads and coordinator before reaching the optimal solution, but the only information needing exchange was the price and quantity of the service. The supply function bidding method was able to clear the market in one shot, but each responsive load must bid a supply curve, which is a function representing the load's willingness to provide different amounts of services at different prices.
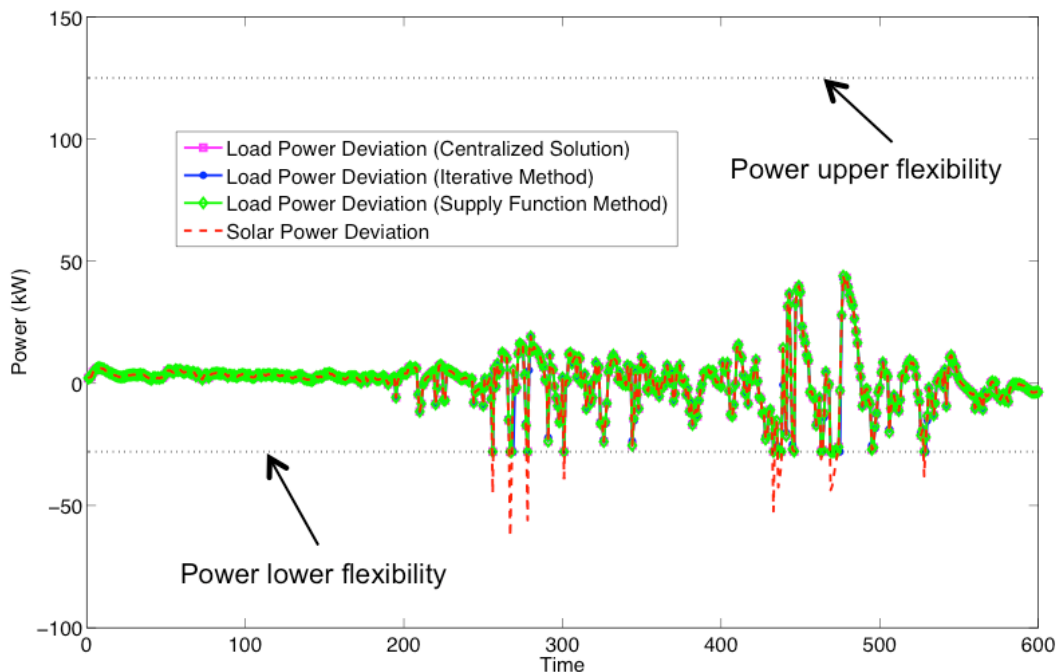

Figure S.4. Solar generation deviation tracking performance using different methods.

AHU Supply Fans

We studied controlling the supply fan in an HVAC system (as shown in Figure S.5) to track a dispatch signal from the coordinator. We developed the fan power models, and identified their model parameters using measurement data from an HVAC system. We proposed two different strategies to control the supply fan, and described the experiment setup and presented the experiment results. We showed that our control methods were able to command the supply fan to track a desired power trajectory, while having little impact on zone thermal comfort.
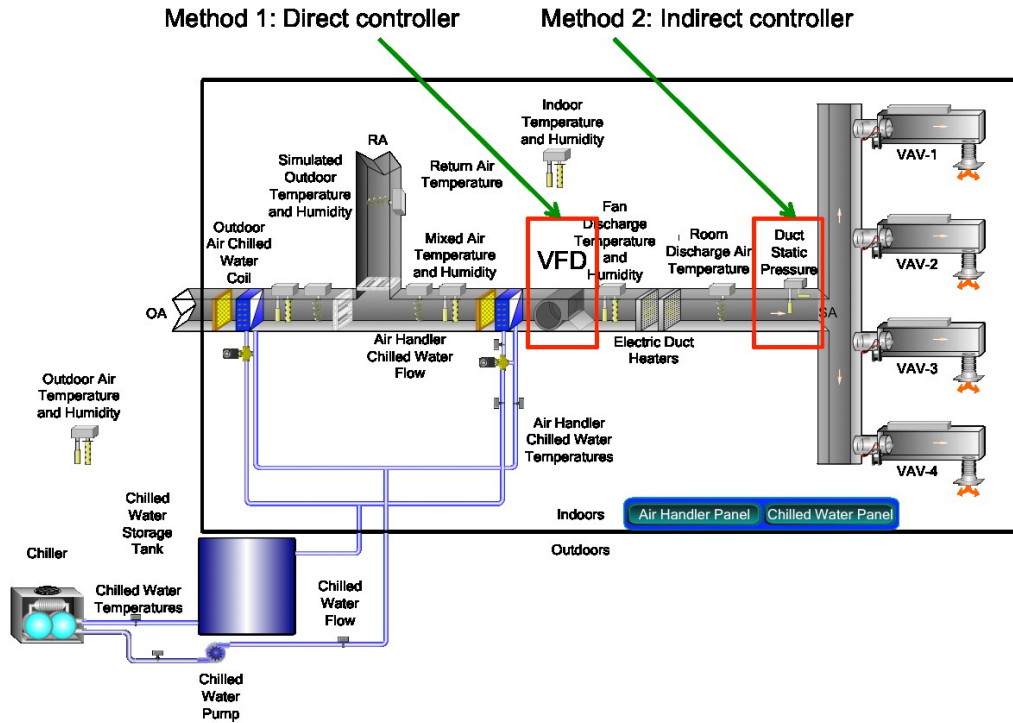
Figure S.5. Schematic of a multi-zone commercial building HVAC system.

To develop the models and identify their parameters, we used measurement data from a commercial building HVAC system on the PNNL campus. The data were collected by VOLTTRON™ and stored in a trended database. A third-order polynomial model was used in the fan power curve. The model parameters were identified via regression and the R-squared was 0.9835. The figure of comparing model calculation and measured power (Figure S.6) showed that the fan power prediction by the model matches the fan power measurement really well.
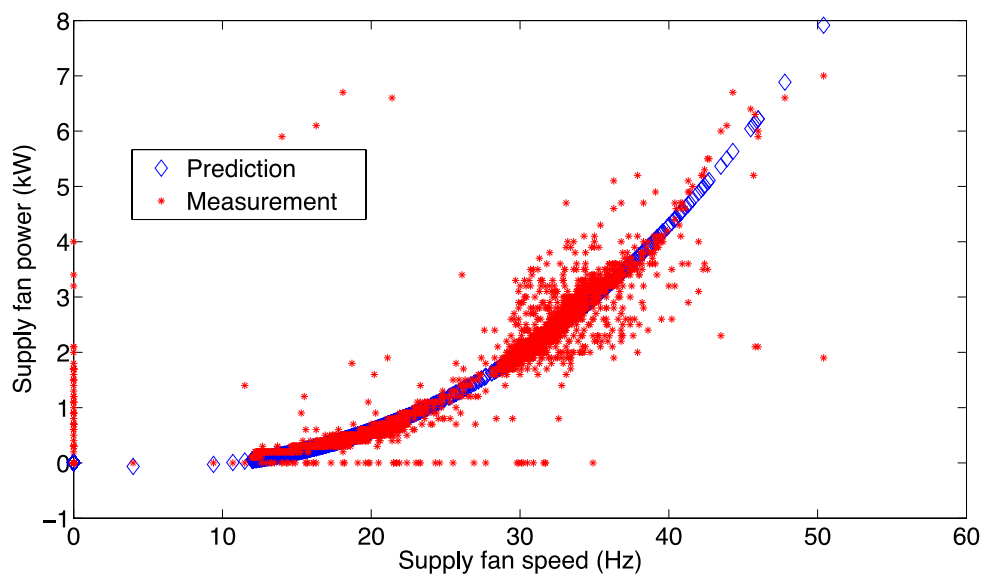

Figure S.6. Prediction and measurement of fan power as a function of fan speed.

In this study, two methods (Figure S.5) were developed to control the fan power. In method (1) fan power was directly controlled by the fan speed, and in method (2) fan power was controlled by the duct static pressure reset (indirect control). In the first method, the variable frequency drive (VFD) speed was directly controlled so that the fan power could reach the desired value. In the second method, a model-free feedback proportional-integral (PI) controller was used, which measured the current fan power and adjusted the duct static pressure so that the fan power could meet the desired value. Method (1) exhibited fast response and accurate tracking performance, but it broke the existing supply fan control loop, in which the fan speed was controlled to maintain a constant duct static pressure. Method (2) did not need to break the control loop, but it required adding another feedback control loop to reset the duct static pressure on top of the pressure controller. Hence, it had certain settling time and transient response (Figure S.7) before the fan power accurately tracked a desired value due to the multiple feedback control loops. Additionally, when tracking a fast changing signal, the tracking performance of the indirect controller was less accurate than that of the first method. Both methods worked fine for the studied HVAC system in which baseline fan power is almost a constant or changes slowly during the occupied hour.
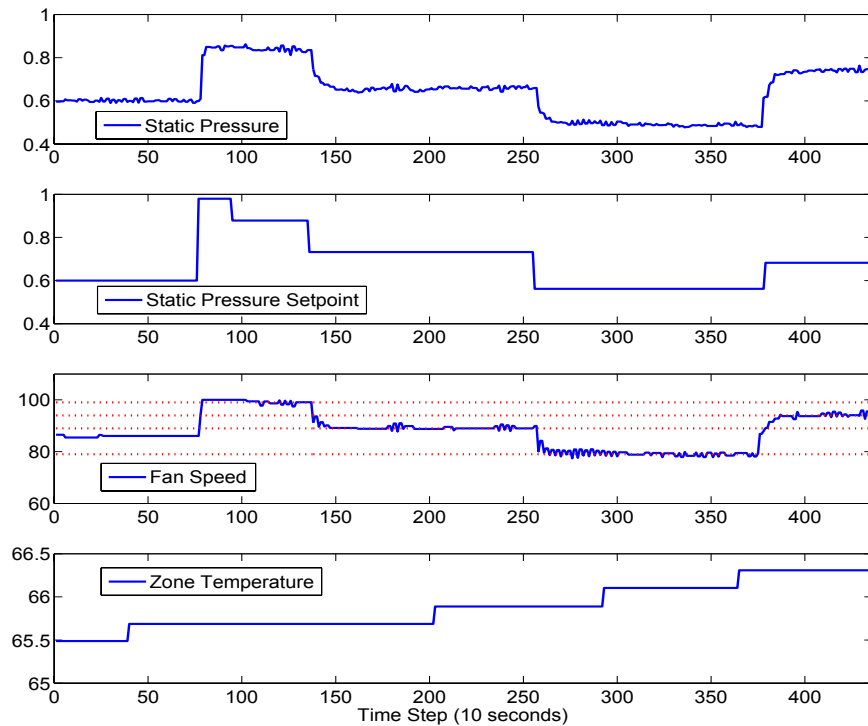

Figure S.7. Experiment result of using the indirect controller.

Packaged Rooftop Units

RTUs are used in about 50% of all commercial buildings, and they serve over 60% of the commercial building floor space in the United States. The configuration of a typical RTU is depicted in Figure S.8. Its main components include the supply fan, heating coil, cooling coil, and outdoor-air, return-air, and relief-air dampers. In our research, we studied the modeling and control of a RTU supply fan for renewable integration. We first described the RTU system and discussed the system modeling and identification. We also described its operating mode and legacy RTU control logic. We next presented the control strategies,

test bed setup, and experiment results. In the end, we summarized the lessons learned from this experiment.
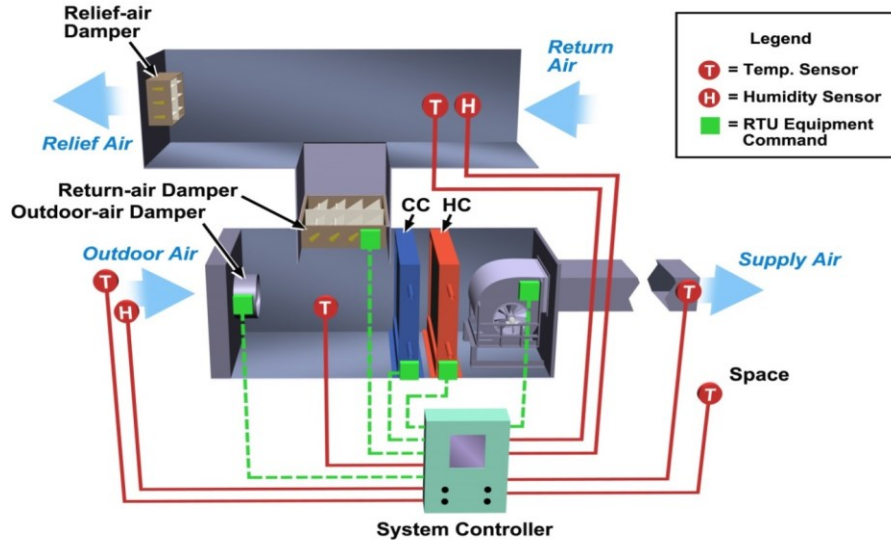

Figure S.8. Schematic diagram of a packaged roof top unit.

To demonstrate control of an RTU to track a dispatch signal from the coordinator, a RTU simulation test bed was set up using Modelica and Dymola. Modelica is an object-oriented modeling language. The developed RTU model was compiled and simulated in Dymola, the execution environment for Modelica. Figure S.9 shows the diagram of the top-level RTU model, which consists of sub-models for physical components of the RTU (such as the heating/cooling source, fan, mixing box), a mode controller, and an HVAC zone.
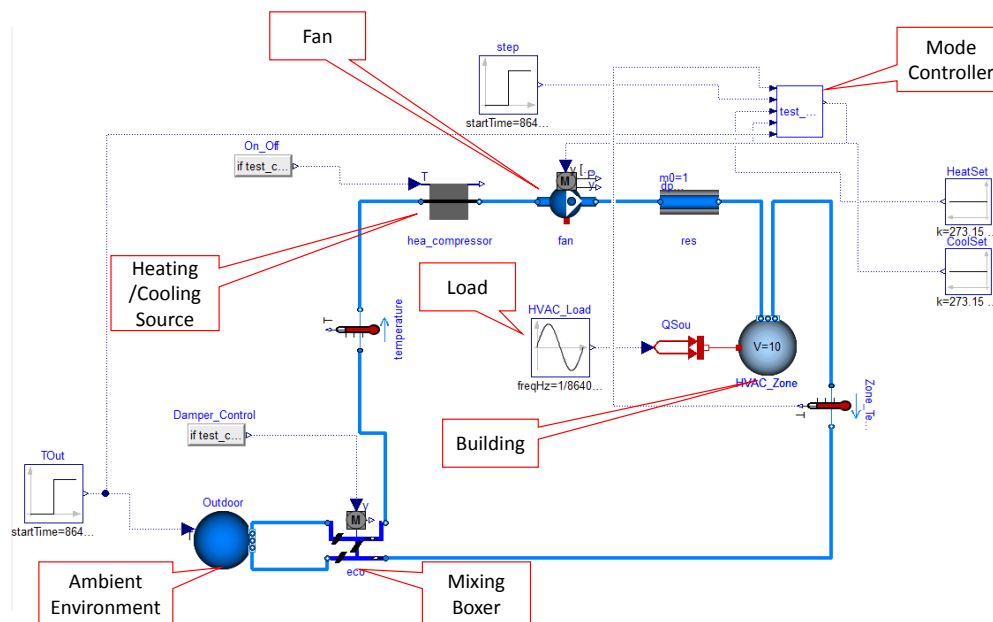

Figure S.9. Diagram of the RTU simulation test bed.

Simulation results (shown in Figure S.10) showed that the fan power deviation could track the dispatch power command very well and the zone temperature had no significant deviations from the baseline case. The results showed that RTUs have a great potential for tracking a power dispatch signal from the grid, while having little impact on the thermal comfort of the served zone. In the future, we will implement the designed control strategies using VOLTTRON on a physical RTU test bed, and examine the potential of tracking a dispatch power signal and its impact on RTU operation and thermal comfort of occupants.
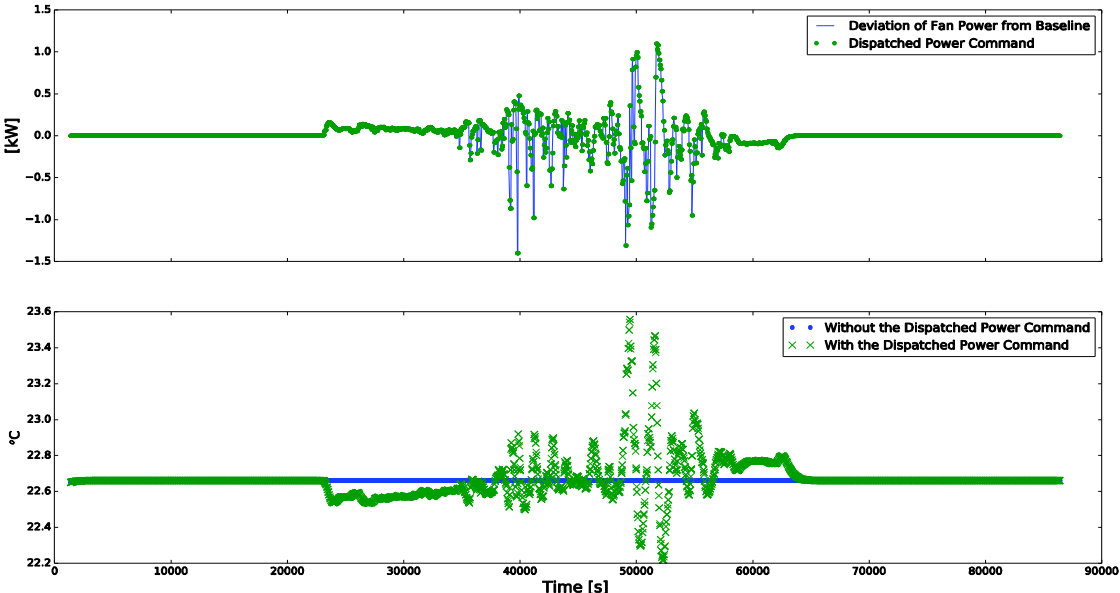

Figure S.10. Simulation results of controlling RTU supply fan for renewable integration.
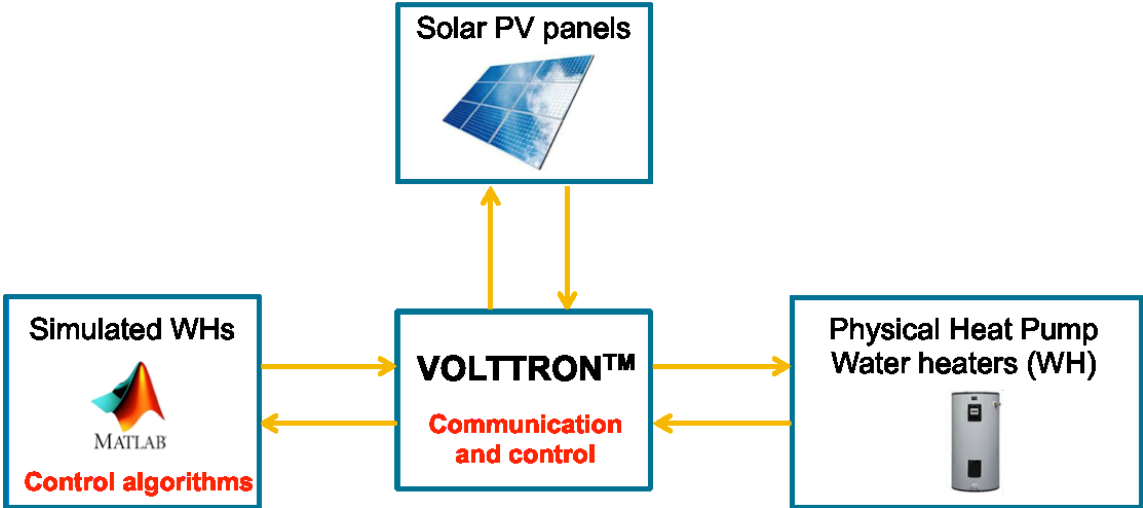
Water Heaters

Figure S.11. Schematic of the hardware-in-the-loop demonstration of coordination of water heaters for renewable integration.

TCLs such as WHs, air-conditioners, heat pumps, and refrigerators, have great potential for assisting renewable integration because of their inherent thermal storage capacities. In our study, we conducted a hardware-in-the-loop (HIL) demonstration of the control of a large aggregated population of WHs to track solar generation deviation. The schematic of the HIL demonstration is provided in Figure S.11. In practice, each WH's power consumption is discrete; it is either the rated power or zero. To get a "smooth" response, a total amount of 1000 WHs were used in the study where 999 of the 1000 WHs were simulated in MATLAB and 1 of the WHs was a physical WH on the PNNL campus (as shown in Figure S.12). The WH models and control algorithms reside in MATALB, and the VOLTTRON™ platform collected data from the solar PV panel and the physical WH, and communicated with MATLAB. Moreover, the VOLTTRON™ platform sent ON-OFF commands to the physical WH.

Figure S.12. Picture of physical WHs.

A centralized control scheme to aggregate and control WHs was used in this study. At each time step, the aggregate power consumption was predicted and compared with the baseline power of the WHs. On the basis of the power difference and the solar generation deviation, a priority-stack-based controller was used to determine which WHs were to be turned ON or OFF. The control commands were then sent to the simulated WHs in MATLAB, and the physical WH via VOLTTRON™. The process was repeated until the end of the considered time horizon.

The experiment was conducted for a 24-hour window. The demonstration results showed that the WHs power deviation from baseline could successfully track the solar generation deviation and the temperatures of simulated and physical WH could be well regulated within the user-specified temperature bounds.
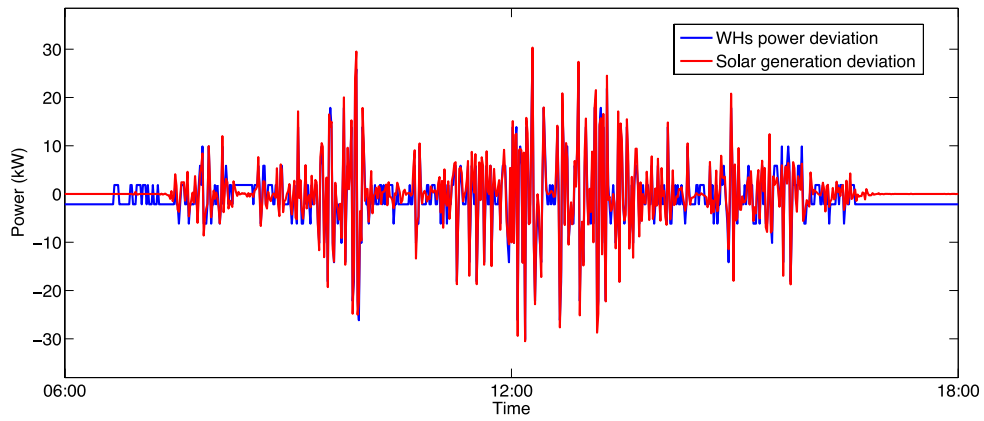
(a) Power tracking performance
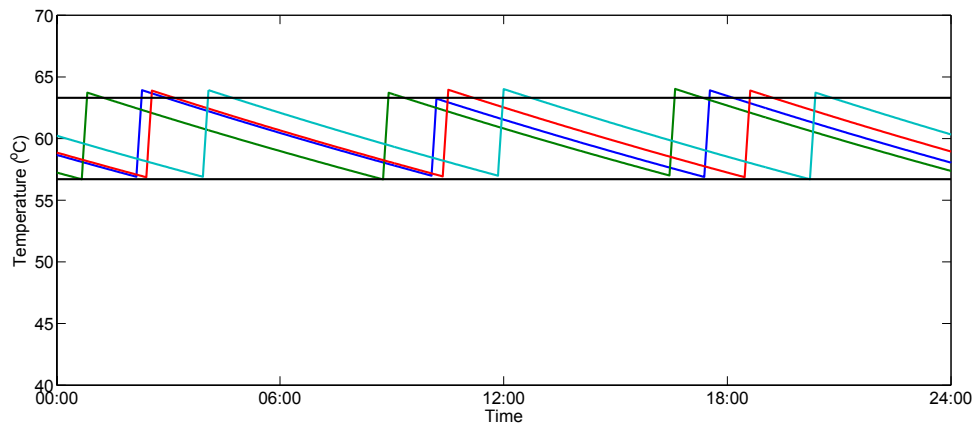
(b) Temperature trajectories of a few simulated WHs.
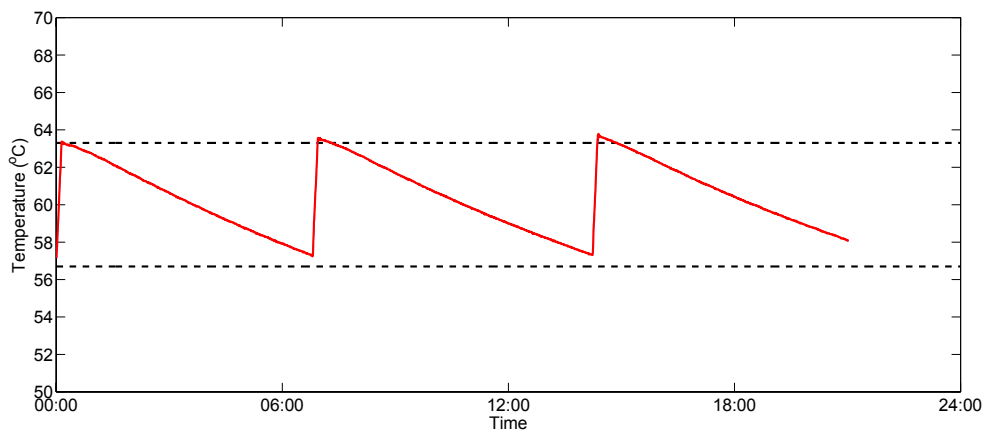
(c) Temperature trajectory of the physical WH.

Figure S.13. Experiment results of coordinating a population of WHs for renewable integration.

In summary, we showed that each of the three types of available loads at PNNL (AHU, RTU, and a collection of WHs) were able to track the power fluctuations of an onsite solar PV panel on the PNNL campus. Moreover, we demonstrated that by providing services to the grid, there was little impact on the thermal comfort of the end-users. In our study, we conducted experiment-based, simulation-based, and HIL-based demonstrations. In the future, we would like to implement the demonstrated control strategies on real test beds and examine their performance and impacts on the quality of service to the end-users. We are also interested in conducting a live demonstration of the coordination of a collection of flexible loads to absorb the power fluctuations of renewable generations.

Acknowledgments

Acronyms and Abbreviations

AHU	air-handling unit
DOE	U.S. Department of Energy
EV	electric vehicle
HIL	hardware-in-the-loop
HVAC	heating, ventilation, and air-conditioning
IP	internet protocol
kW(s)	kilowatt(s)
MCP	market clearing price
MPC	model predictive control
PID	proportional-integral-derivative (controller)
PNNL	Pacific Northwest National Laboratory
ppm	parts per million
PV	photovoltaic
RTU	rooftop unit
TCL	thermostatically controlled load
UK	United Kingdom
VAV	variable air volume
VFD	variable frequency drive
WH	water heater

Contents

Figures

Tables

1.0   Introduction

Environmental concerns have imposed an urgent need for energy system revolution. It is believed that the future energy system will rely substantially on renewable energy resources such as wind and solar power. However, these two variable generations are very uncertain in their productions, and therefore bring substantial power fluctuations to the power system. Buildings, in which we live and work, have large inherent power flexibility, and they are excellent candidates to absorb the uncertainty of variable generations. In this section, we discuss the background and motivation of building-to-grid integration, and present the literature review of coordinating and controlling various types of flexible building loads to provide services to the grid.

1.1   Background and Motivation

Global warming and climate change are perhaps one of the greatest threats to our planet. According to National Aeronautics and Space Administration Goddard Institute of Global Studies, the global mean-surface-temperature has increased by 0.6 °C in the past 40 years (NASA Goddard Institute for Space Studies n.d.). Moreover, the global surface temperature is projected to rise a further 0.3 to 1.7 °C under the Institute's lowest emissions scenario using stringent mitigation and 2.6 to 4.8 °C under its highest scenario at the end of the 21st century (Stocker, et al. 2014). To a great extent, global warming is due to excessive exploitation of fossil fuel energies. Currently, more than 80% of our energy is derived from fossil-fuel resources such as oil, coal, natural gas, etc. (International Energy Agency 2014). If the current energy situation continues, it will exacerbate global warming, and threaten our environment.

Decarbonizing the global energy system is one of the most important endeavors of our time. It is widely accepted that a sustainable energy future will rely substantially on renewable energy sources such as wind and solar. However, variable generation sources are uncertain, intermittent, and stochastic. Large-scale integration of variable generations will inject substantial power fluctuations into the power gird, which makes maintaining the power balance between supply and demand very challenging. Currently, the variability of renewables is handled through supply-side reserves. However, recent studies show that deep penetration of renewables will substantially increase the need for generation reserves (Smith, et al. 2007), (Makarov, et al. 2009), (Helman 2010). If additional reserves are provided by fossil-fuel generators, the net carbon benefit from renewables will be diminished, generator efficiency will reduce, and an economically untenable generation mix will result. Hence, the stability and reliability of the grid will require more flexible consumption from demand-side resources.

Flexible building loads, such as commercial building heating, ventilation, and air-conditioning (HVAC) systems, residential thermostatically controlled loads (TCLs), electric vehicles (EVs), and deferrable appliance loads, have a great potential to provide various services to the grid. For example, by slightly changing the ON/OFF schedule of a collection of WHs, their aggregate power deviation from baseline can be made to follow the power deviation of a solar photovoltaic (PV) panel, thereby mitigating the power output variability of the PV panel.

This project aims to investigate how to efficiently and economically integrate renewable energies into the grid by aggregating, coordinating, and controlling distributed demand-side resources. Intelligent control of these resources will help the grid absorb the uncertainties from renewables without procuring additional expensive and polluting generation reserves. Our research develops control algorithms and

testbed infrastructure to demonstrate coordination and control of flexible building loads for renewable integration. We also conduct experiments of controlling various types of flexible building loads to make net load (building load minus renewable generation) as smooth as possible so that fewer power fluctuations are injected into the power grid.

1.2 Literature Review and Research Objective

Generally speaking, flexible loads can be placed into three broad categories: commercial building loads, residential building loads, and other types of loads. In this section, we present a comprehensive catalog of flexible loads and review the literature study of each type of resources. Table 1.1 summarizes the detailed list of flexible loads by building category.

Table 1.1. A catalog of flexible building loads[1].

Commercial Building Loads					
Load Type	**Control Strategies**	**Response Speed**	**Response granularity**	**Response Duration**	**Rated Power**
VFD-controlled supply fans in air-handling units (AHUs)	Control fan speed, total supply air flow rate, or duct static pressure	Seconds to minutes	Continuous	10-30 minutes in either direction from baseline	Varies, typically 10-100 kW
VFD-controlled rooftop units (RTUs)	Control fan and compressor speed, turn ON/OFF, or change temperature set point	Seconds to minutes	Continuous	10-30 minutes in either direction from baseline	Varies, typically 10-100 kW
Staged RTUs	Set at different stage, turn ON/OFF, or change temperature set point	Seconds to minutes	Discrete	10-30 minutes in either direction from baseline	Varies, typically 10-100 kW
Chiller	Change supply chilled water temperature set point, control fan, or compressor	Seconds to minutes	Continuous or discrete depending on its type	30-60 minutes in either direction from baseline	Varies, typically 10-1000 kW
Lighting	Stepped or continuous dimming, or switch OFF	Seconds	Continuous and discrete	Infinite time in either direction from baseline	25% of a commercial building total electrical power (Building Energy Data Book n.d.)

---

[1] The numbers in this table are mainly based on the authors' experience and heuristics.

Plug loads (desktops, table lambs)	Turn ON/OFF	Seconds	Discrete	Minutes to hours	See (Weightman and Field 2012)
Residential Building Loads					
Thermostatically controlled loads (air-conditioners, heat pumps, WHs, refrigerators)	Turn ON/OFF or change temperature set point	Seconds to minutes	Discrete	10 minutes to hours in either one direction from baseline	
Distributed energy storage (e.g., Tesla Powerwall, flywheels)	Charge and discharge	Seconds	Continuous	Seconds to hours	
Deferrable appliances (pool pumps, dryer, etc.)	Turn ON/OFF	Seconds	Discrete	Minutes to hours	
Electric vehicles	Turn charging ON/OFF	Seconds	Continuous and discrete	Minutes to hours	
Plug loads (desktops, laptops, table lambs)	Turn ON/OFF	Seconds	Discrete	Minutes to hours	
Other Loads					
Refrigeration Systems		Seconds to minutes		hours	
Data centers	Several control strategies are given in (Ghatikar and al. 2012), (Wierman, et al. 2014)	Seconds to minutes	Continuous and discrete	Minutes to hours	100s kW
Aluminum manufactories	Control aluminum smelter	Seconds	Continuous	Minutes to hours	
Desalination plants					
Municipal watering systems					
Agriculture irrigation systems	Control VFD pumps	Seconds	Continuous	Minutes to hours	

Energy Storage	Charge and discharge	Seconds	Continuous	Minutes to hours	Varies, see (DOE Global Energy Storage Database n.d.)

Coordination of flexible loads for various grid services has been extensively studied in the literature, and various modeling and control strategies have been proposed. For example, control of the supply fan in a commercial building HVAC system to provide frequency regulation service was studied by (Hao, Lin, et al. 2014), (Y. Lin, P. Barooah and S. Meyn, et al. 2015), (Vrettos, Kara, et al. 2016), (Vrettos, Kara, et al. 2016). (Lin, Barooah and Meyn 2013), (Lin, Barooah and Mathieu 2015) investigated using the combination of a fan and a chiller in a commercial HVAC system to provide slower ancillary services. Demonstration of the coordination of four RTUs to track solar generation was presented by (Nutaro, et al. 2015). Various control strategies and several experiment demonstrations using commercial buildings for grid services were reported in (Motegi, et al. 2007), (Kiliccote, et al. 2010), (Yin, et al. 2010). In (Mathieu, Koch and Callaway 2013), (Callaway 2009), (Hao, Sanandaji, et al. 2015), (Hao, Sanandaji, et al. 2013), (Mathieu, Kamgarpour, et al. 2013), (Malhame and Chong 1985), (Bashash and Fathy 2013), (Zhang, et al. 2013), (Ihara and Schweppe 1981), (Kundu, et al. 2011), (Lu and Chassin 2004) reported on the conduct of substantial research on aggregation and control of a large population of residential TCLs such as air-conditioners, heat pumps, WHs, and refrigerator for various grid services (e.g., frequency regulation, energy arbitrage, peak load management, etc.). The Markov decision process and mean-filed game were used to control residential pool pumps for ancillary services (S. P. Meyn, et al. 2015), (S. Meyn, et al. 2013). Characterizing flexibility and real-time scheduling of residential deferrable loads for renewable integration has been studied in (Nayyar, et al. 2013), (Subramanian, et al. 2013), (Hao and Chen 2014), (Chen, Ji and Tong 2012), (Papavasiliou and Oren 2010). (Chen, et al. 2010), (Li, Chen and Low 2011), (Li, et al. 2016) developed market-based distributed coordination mechanisms for various types of residential loads. Manufacturing companies and agriculture farms have also been successfully engaged by ramping up and down their power consumption in response to the requirements of the grid (Todd, et al. 2009), (Department of Energy n.d.). Other types of loads such as lighting, plug loads, refrigeration systems, data centers, desalination plants, municipal watering systems, etc. were also considered for demand response and providing ancillary services (Hirsch, et al. 2015), (California Energy Commission's Public Interest Energy Research Program 2011), (Weightman and Field 2012), (Ghatikar and al. 2012), (Wierman, et al. 2014), (Kunczynski and Burger 2014), (Brzozowski 2011). The authors believe that there are also many exceptional works that are not referenced, but conducting a more comprehensive literature review is beyond the scope of this report.

Most of the current studies are limited to modeling, aggregation, and control of a single type of building load, and ignore coordination with other types of flexible resources. To fully exploit the potential of flexible loads, it is essential to study optimal coordination of multiple flexible loads and consider their diversified response characteristics and availabilities. Most of the research to date focuses on proof-of-concept demonstrations by conducting simulation-based studies. But to support large-scale applications, it is important to demonstrate the feasibility of the proposed approaches by conducting experiments on physical devices and learn lessons from field tests. In addition, many existing experiment demonstrations have used device-specific test bed and software tools to communicate with and control the physical loads. A low-cost and open-source platform to manage a diverse collection of flexible loads does not exist.

To address the above issues, we study coordination of a collection of heterogeneous flexible loads. The use case we focus on in this report is renewable integration, which aims to use flexible building loads to locally absorb the uncertainty of renewables so that fewer power fluctuations are injected into the power grid. Based on the available resources, we pick three types of flexible building loads: 1) a supply fan in an AHU, 2) a supply fan in an RTU, and 3) a collection of WHs. Additionally, we demonstrate coordination and control of these three types of building loads to compensate for the solar generation deviation of a solar PV array on the PNNL campus. The overall objective is to coordinate and control the buildings' power consumption to mitigate the power fluctuation of the solar PV panel without affecting the quality of service to the end-users. Moreover, our test bed infrastructure is based on VOLTTRON™ (Pacific Northwest National Laboratory 2016)—a PNNL-developed open source software platform that facilitates communicating with and controlling flexible building loads. It is worth mentioning that although our focus is on the three specific types of building loads and to achieve renewable integration, the developed coordination and control algorithms and the test bed infrastructure can be adopted to provide other building-to-grid functions.

1.3 Report Organization

The rest of this report is organized as follows. Section 2 presents solar data analysis for the solar PV array on the PNNL campus. In Section 3, we propose a coordination framework for the supply fan in the AHU, the supply fan in the RTU, and WHs. The control algorithms, test bed configurations, and experiment results for the three different types of building loads are described in Sections 4, 5, and 6, respectively. The report ends with conclusions and a discussion of future work in Section 7.

2.0   Solar Data Analysis

In this section, we describe the solar data source used in our research and identify its characteristics. Welch's method is used to estimate the power density of the solar generation data at different frequencies. We find that the solar PV power outputs are very intermittent, and the uncertainty occurs at most of the frequencies. We also examine the solar generation data from other sources, and find similar solar generation characteristics. Moreover, we present the methods we used for solar power prediction in our research, and state the objective of our research.

2.1   Data Source

The solar PV array (Figure 2.1) is located in Richland, Washington, which has about 200 sunny days a year. Richland receives about 7 inches of rain and 8 inches of snowfall per year, while the annual national average of rain and snowfall are 37 inches and 25 inches, respectively. This makes Richland a great location for installing solar PV arrays. The solar data we use are obtained from the solar PV array shown in Figure 2.1. Its maximum power production rate is 126 kW. It produces about 175,000 kWh of electricity in a year and provides power to the Environmental Molecular Science Laboratory (EMSL), which includes a super-computing facility and adjacent EV-charging stations. We use the VOLTTRON^{TM} software platform to sample its power production data at a time resolution of 1 minute, and store the data in a Mongo database using the VOLTTRON^{TM} historian agent.

Figure 2.1. Solar PV panels on PNNL campus [source: Hire Electric]

2.2   Data Analysis

The summers in Richland, Washington, are hot and thunderstorms are infrequent; winters are mild and falls snow only occasionally. To examine the characteristics of the solar generation data during different seasons, we selected the historic solar PV output for two different time periods. Figure 2.2 (a) and (b) show the solar PV generation from October 23, 2015 to November 9, 2015, and from March 15, 2016 to April 6, 2016, respectively. We observe that the solar output has a typical diurnal pattern, while the period

from October 23, 2015 to November 9, 2015 (late fall) has more cloudy days than the period from March 15, 2016 to April 6, 2016 (early spring). Additionally, Figure 2.3 (a) and (b) depict the daily solar PV output during the two different time periods. In particular, different colors in the 2.s represent different days. As can be seen in the figures, the solar PV panels produce at most 108 kW of power, which is less than the designed capacity of 126 kW. Additionally, the number of production hours during the period from October 23, 2015 to November 9, 2015 (late fall) is less than that during the period from March 15, 2016 to April 6, 2016 (early spring). More importantly, the solar output is very intermittent. On most of the days, its power output fluctuates frequently and abruptly due to passing clouds.
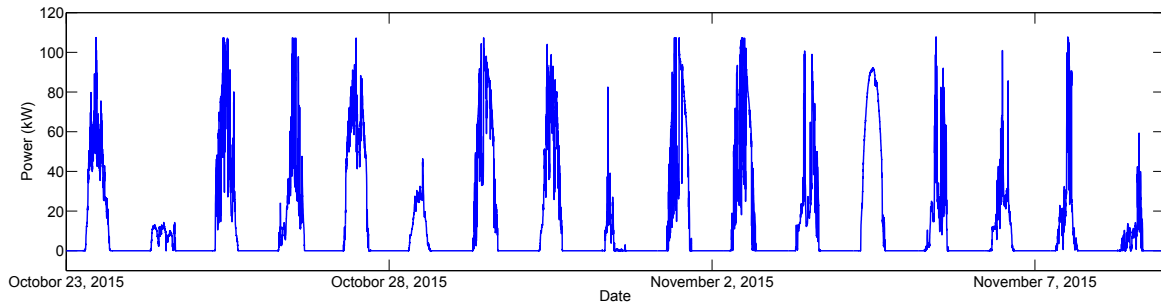
(a) Solar PV output from October 23, 2015 to November 9, 2015 (17 days).
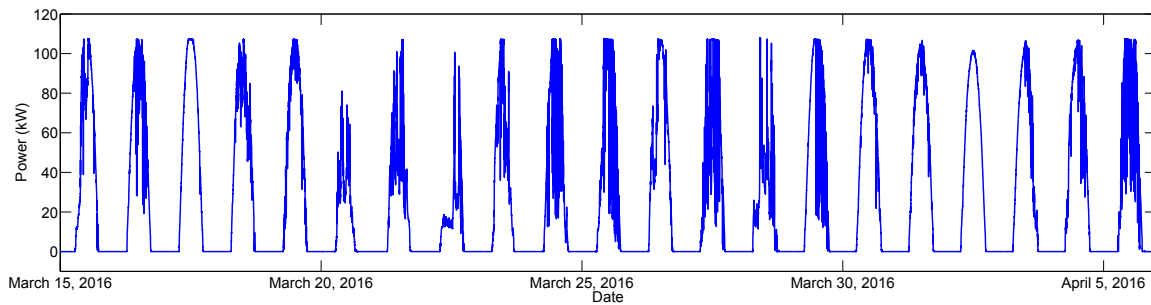
(b) Solar PV output from March 15, 2016 to April 6, 2016 (21 days).

Figure 2.2. Historic solar PV generation data.
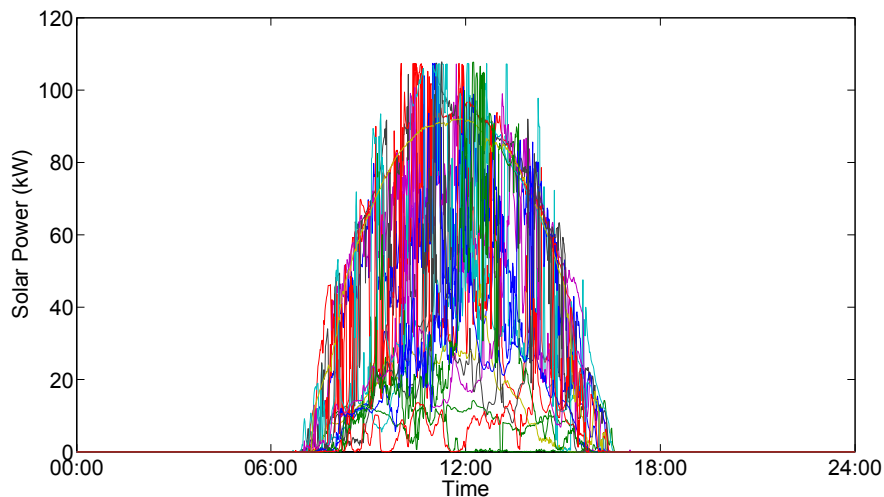
(a) Daily solar PV outputs from October 23, 2015 to November 9, 2015. Different color represents different days.

2.2
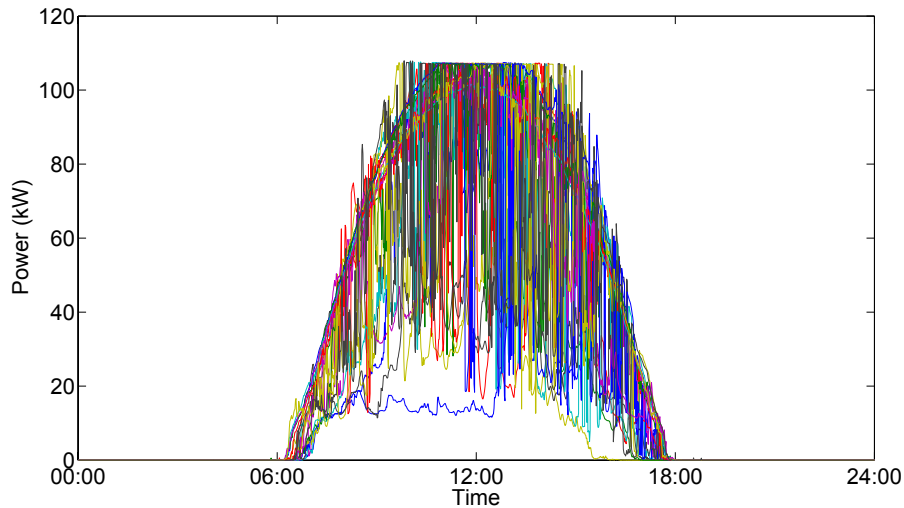
(b) Daily solar PV outputs March 15, 2016 to April 6, 2016. Different color represents different days.

Figure 2.3. Daily solar generation data.

We used Welch's method to estimate the power density of the solar generation data at different frequencies. Figure 2.4 (a) and (b) show the power spectral density estimations for the solar data in the two aforementioned time periods. The power spectral density provides a useful way to characterize the amplitude versus frequency content of a signal. We observe in Figure 2.4 that the solar generation has large density in the low-frequency band [0, 0.001]. This is mainly because the solar generation exhibits diurnal behavior. Additionally, it has relative constant density in the high-frequency band [0.001, $1/(2\times60)$] Hz, where $1/(2\times60)$ is the highest detectable frequency because the sampling rate is every 60 seconds (1 minute). This implies that the uncertainty of solar generation occurs at most high frequencies, and the uncertainty occurs at most of the frequencies, and therefore leaves no choice to target on some particular frequency band as suggested by the approach of (Barooah, Buic and Meyn 2015), (Hao, Lin, et al. 2014), (Lin, Barooah and Meyn 2013).
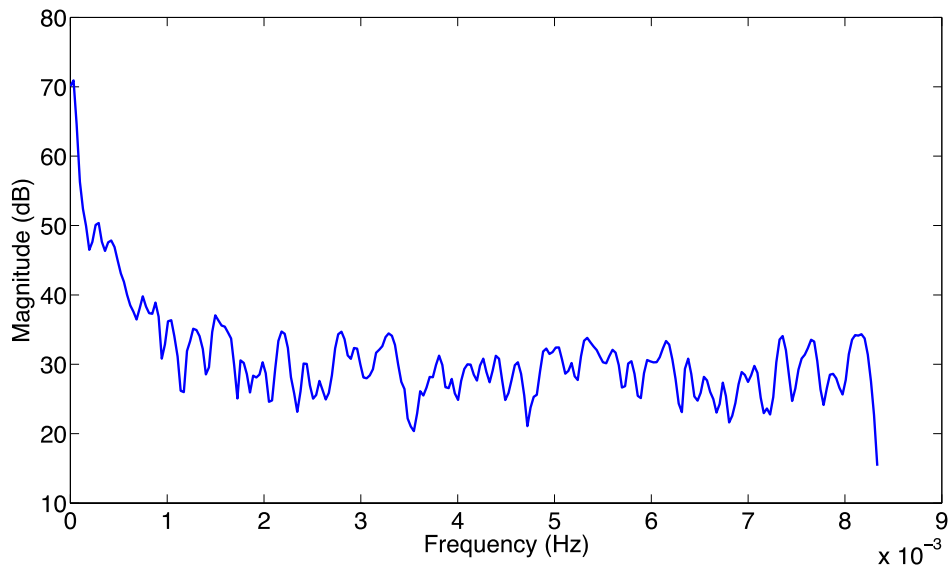
(a) Power spectral density of the solar generation data from October 23, 2015 to November 9, 2015.

2.3
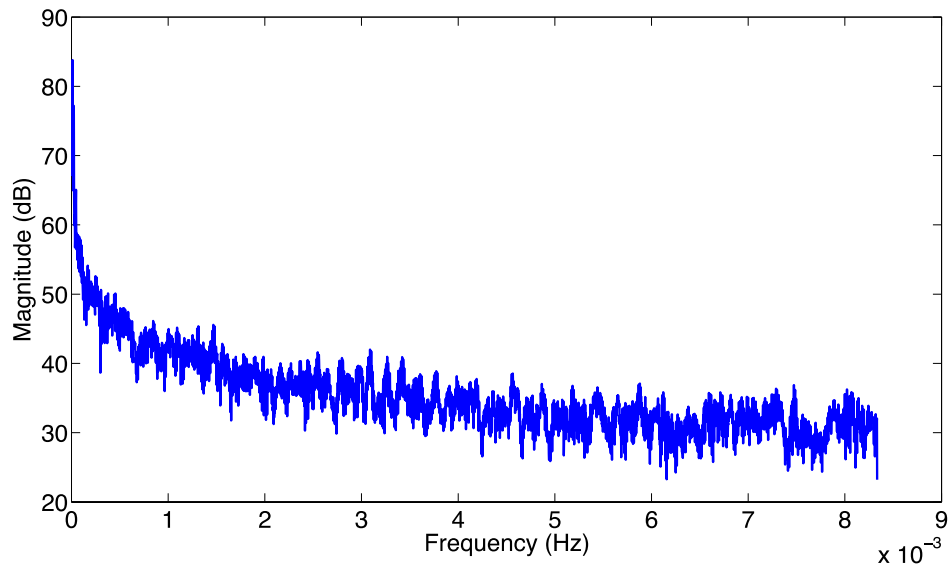
(a) Power spectral density of the solar generation data from March 15, 2016 to April 6, 2016.

Figure 2.4. Welch's power spectral density estimate of the solar generation data.

2.3   Solar Prediction

Solar power forecasting is a very active research area, and it is very important for planning the operation of power systems. Many different methods with different granularities, accuracies, and complexities have been proposed; see (Bacher, Madsen and Nielsen 2009), (Chen, et al. 2011), (Lorenz, et al. 2009) for a few examples.
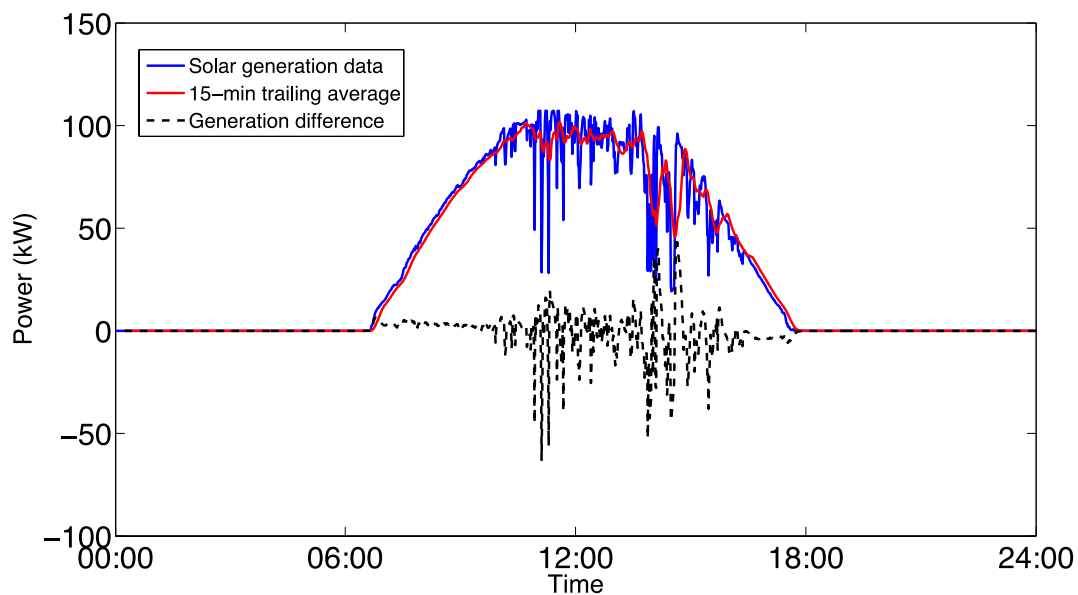
Figure 2.5. Daily solar generation, its 15-minute trailing average, and the difference between the two for solar data on March 15, 2016.

The objective of this research is to coordinate and control flexible building loads to compensate for the uncertainty of solar PV generations. Solar power prediction will help us to determine how power used by building loads should be adjusted up or down so that the loads' real-time power consumption minus their baseline power tracks the difference between the real-time solar power production and its prediction. However, developing an accurate solar prediction method is far beyond the scope of our current research. For simplicity, we use the 15-minute trailing average of the solar generation as the solar prediction in the next time step. Mathematically, at each time step t, we predict the solar PV output at the next time step using the following formula:

$$\hat{P}_{t+1}^{S} = \frac{\sum_{k=t-14}^{t} P_{k}^{S}}{15},$$

(1)

where P_{t}^{S} is the solar power measurement at time step t. Figure 2.5 depicts the solar generation, its 15-minute trailing average, and the difference between the two for the solar output measured on March 15, 2016. We observe that the 15-minute trailing average is much smoother than the ground-truth solar power output.

Another option is to take ideal solar generation on a clear-sky day as the solar power prediction. For example, we can see from Figure 2.2 (b) that the solar generation on March 17, 2016 is an ideal generation case. Figure 2.6 shows the solar generation (March 15, 2016), ideal generation (March 17, 2016), and their difference between the two. Moreover, if other solar power prediction methods are used, we can replace the red lines in Figure 2.5 or Figure 2.6 with the output of those methods, and recalculate the generation deviation (the black dashed lines in Figure 2.5 or Figure 2.6). The objective of our research is to control the building loads so that their real-time power minus their baseline tracks the solar generation deviation.
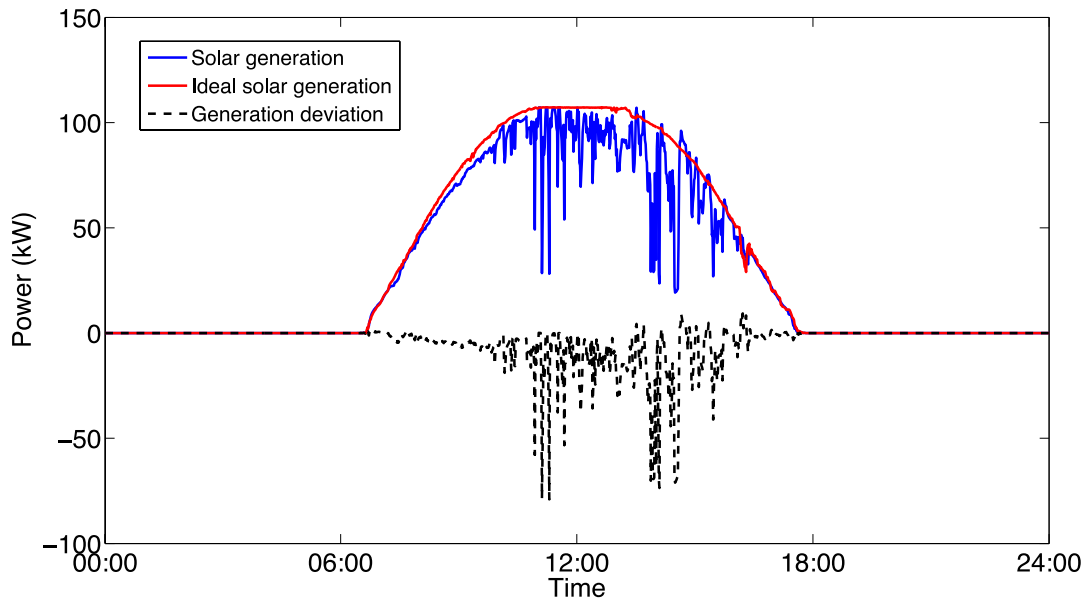
Figure 2.6. Daily solar generation (March 15, 2016), ideal solar generation (March 17, 2016), and their difference between the two.

2.5

The characteristics of the solar generation we identified in this report are not limited to the solar PV array on the PNNL campus. We have also checked the data from other sources. For example, Figure 2.7 depicts the selected daily solar PV generation data from the Newquay Weather Station in the United Kingdom (UK) (Newquay Weather Station n.d.). It is found that the characteristics of the solar generation output from solar PV panels measured in the UK are consistent with what we have observed.
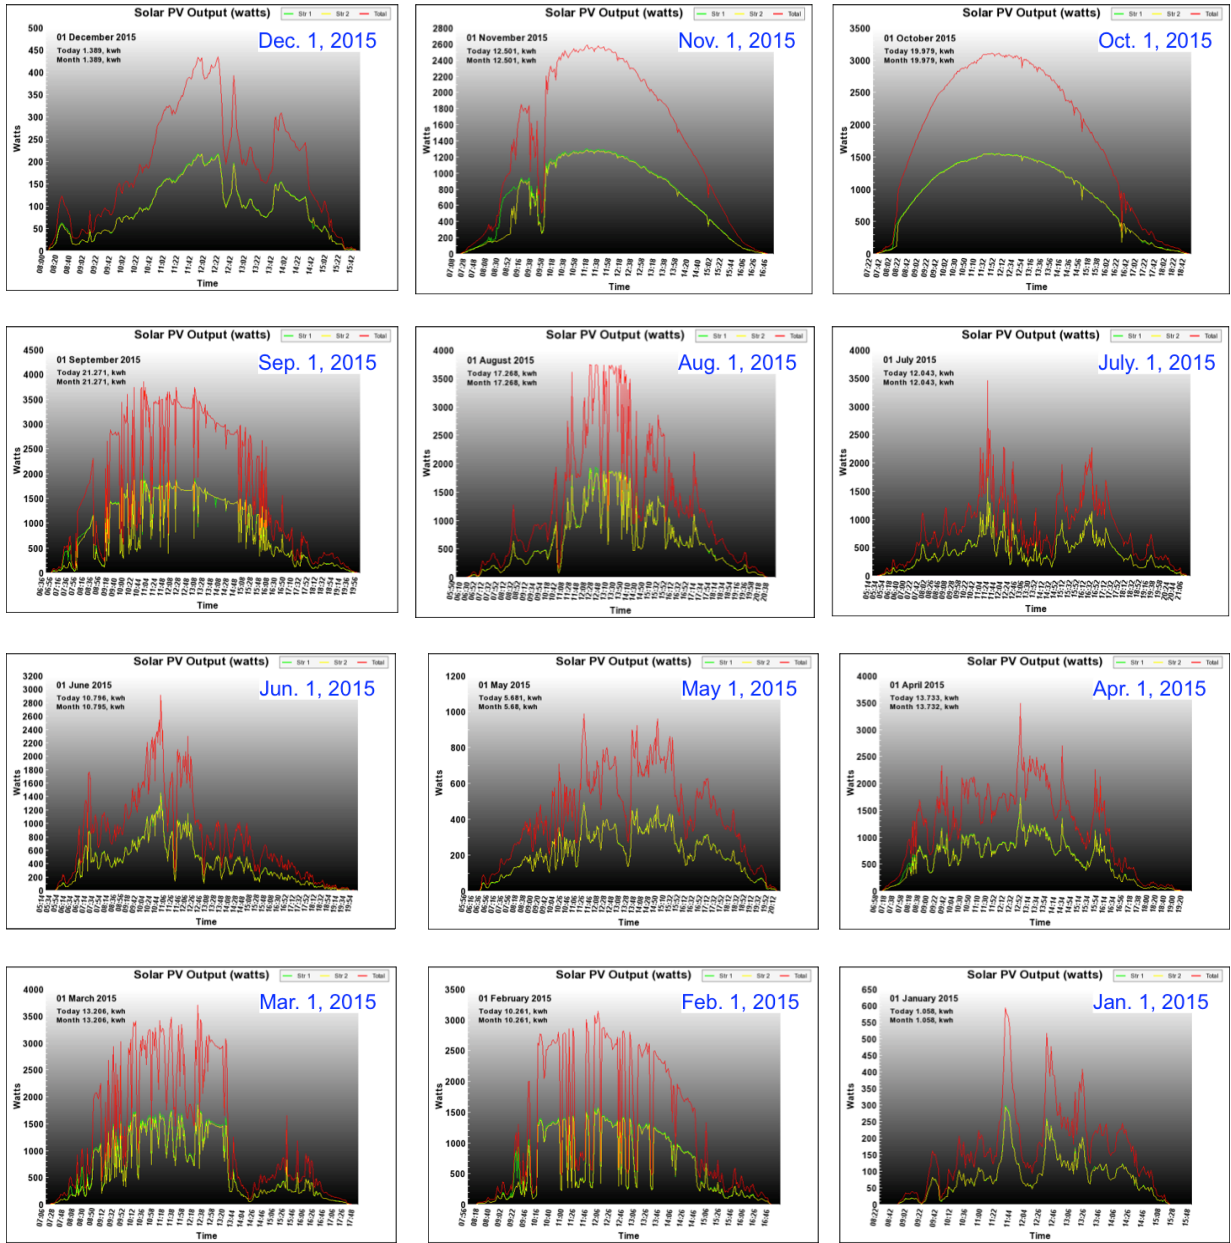
Figure 2.7. Solar generation of solar PV panels measured in the UK for the first day of each month in 2015.

3.0 Transactive Coordination of Flexible Building Loads

In this section, we present a two-level coordination framework to manage flexible building loads for renewable integration. We propose two transactive coordination strategies to engage self-interested responsive loads to provide services to the grid. Simulation studies are also conducted to demonstrate the efficacy of the coordination framework.

3.1 Coordination Framework

To fully exploit the flexibilities of various types of loads, we proposed to use a coordination framework to manage their power consumption. A schematic representation of the coordination framework is depicted in Figure 3.1. The framework consists of two levels: the flexible load level and the coordinator level. The lower level includes different types of flexible loads. At each time step, these loads report their forecasted baseline for the next period and their power flexibility (the range its power is able to change from the baseline). In the upper level, there is a coordinator who is responsible for forecasting the renewable generation and implementing coordination strategies to manage the flexible loads so that their real-time power minus the baseline power can track the renewable generation deviation. In this report, we focus on three types of building loads for demonstrations—a VFD-controlled fan in an AHU, a VFD-controlled fan in an RTU, and a population of WHs. These three types of loads can be categorized into continuous loads (supply fans in the AHU and RTU) and discrete loads (WHs). For continuous loads, we can modulate their power consumption continuously to track a dispatched power trajectory. However, this is not the case for discrete loads, because the power of each WH is either zero (OFF) or the rated power (ON). Therefore, we need to aggregate discrete loads and represent them as a single resource to the coordinator, so that the loads can provide a smoother aggregate response.
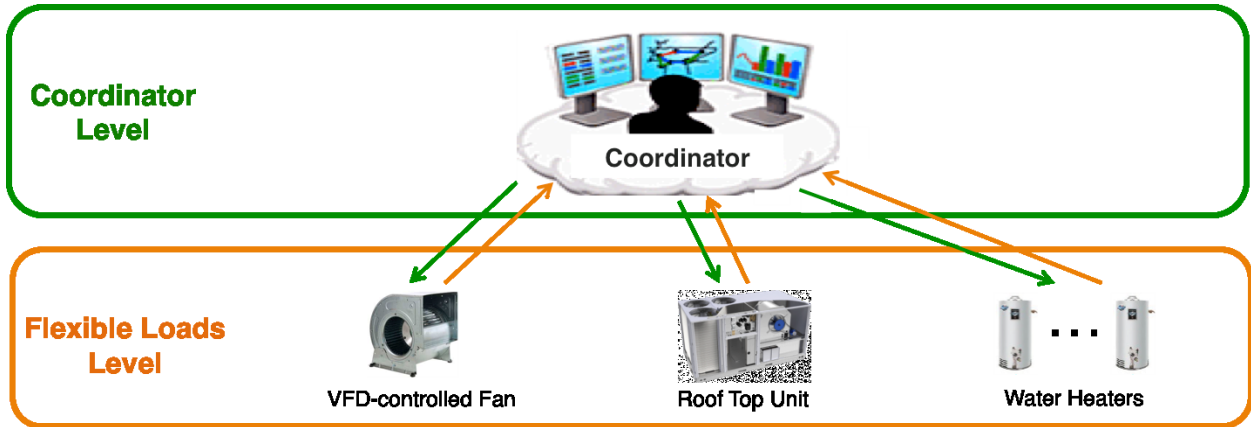
Figure 3.1. Schematic of the hierarchical coordination framework.

3.2 Transactive Coordination Strategies

In our research, we are interested in the transactive control of flexible building loads. Transactive control is a type of distributed control strategy that uses market mechanisms to engage self-interested responsive loads to provide services to the grid. The core of transactive control is the concept of supply-demand market equilibrium.

In our research, the problem of the transactive coordination of flexible building loads for renewable integration is formulated as follows. At each time step, the coordinator forecasts and measures the renewable generation and determines the amount of power that the loads need to reduce or increase from their baseline. For each continuous load or a collection of discrete loads, a power reduction or increase will incur a cost, which is represented by $C_i(p_i)$, where i is the index of that continuous load or collection of discrete loads, and p_i is its power reduction or increase. The control objective is to minimize the total cost subject to the constraint that the total load reduction or increase is equal to the renewable generation deviation, Q. Mathematically, the problem is formulated as follows:

$$\underset{p_i}{\text{minimize}} \quad \sum_i C_i(p_i)$$

subject to:

$$\sum_i p_i = Q,$$

$$-\underline{p_i} \leq p_i \leq \overline{p_i}, \quad\quad\quad (2)$$

where $-\underline{p_i}$ and $\overline{p_i}$ are the lower and upper power limits of the i^{th} load, respectively.

Generally speaking, there are two transactive (or market-based) approaches to solving the above optimization problem: 1) an iterative method based on dual decomposition, and 2) a one-shot supply function bidding method.

3.2.1   Iterative method based on dual decomposition

We first present the iterative method. Upon observing the separability of the cost function and the constraints in Equation (2), we can use the dual decomposition approach (Boyd and Vandenberghe 2004), (Chen, et al. 2010) to develop an iteration-based distributed algorithm to find the optimal solution. More specifically, the optimization problem (Equation (2)) can be decomposed into a master problem and a group of sub-problems. The master problem can be solved by the following gradient-decent method:

$$\lambda^{k+1} = \lambda^k + \alpha(Q - \sum_i p_i^k), \quad\quad\quad (3)$$

where k is the iteration number, λ^k is the price or Lagrange dual variable at iteration k, and α is a small step size. Additionally, for each i, its sub-problem is given by

$$\underset{p_i}{\text{minimize}} \quad C_i(p_i^k) - \lambda^k p_i^k$$

subject to:

$$-\underline{p_i} \leq p_i^k \leq \overline{p_i}. \quad\quad\quad (4)$$

The interpretation of the above iterative algorithm can be explained as follows. First of all, the coordinator broadcasts an initial price λ^0 to the loads and asks them how much power they would like to

reduce or increase. Based on the service price, each load determines the optimal power reduction or increase, p_i, by solving Equation (4), and sends this quantity back to the coordinator. Upon receiving all the p_i^k's, the coordinator calculates if the collective power reduction or increase is equal to the renewable generation deviation Q. If not, the coordinator will update the price using Equation (3), and broadcast it to the loads again. By conducting such an iterative information exchange, the price λ^k and the power reduction or increase p_i^k's will converge to the optimal solutions, λ^* and p_i^*'s.

3.2.2   One-shot supply function bidding method

In the one-shot supply function bidding method, flexible loads provide service to the grid at marginal revenues that are represented by supply curves, and the coordinator buys service with a marginal benefit that is described by a demand curve. If the demand curve intersects the aggregate supply curve, the market is cleared. The price and quantity of the service are determined by their intersection.

In this method, we separate the power reduction and power increase into two different services. In this report, we only present the supply function bidding and market clearing for the power increase service; the power reduction service follows in a similar way. Each continuous load or a collection of discrete loads submits a supply curve (as shown in Figure 3.2 (left)) to the coordinator. Upon receiving all of the supply curves, the coordinator constructs an aggregate supply curve (the red curve in 3.2 (middle)) by summing the individual supply curves. The quantity of service is designated Q, therefore the intersection of the supply curve and the vertical line passing the point $(Q, 0)$ determines the market clearing price λ^*. The coordinator then sends the price λ^* to the loads, and each load determines its power reduction or increase p_i^*, based on the market clearing price λ^* and its own supply curve (Figure 3.2 (right)).
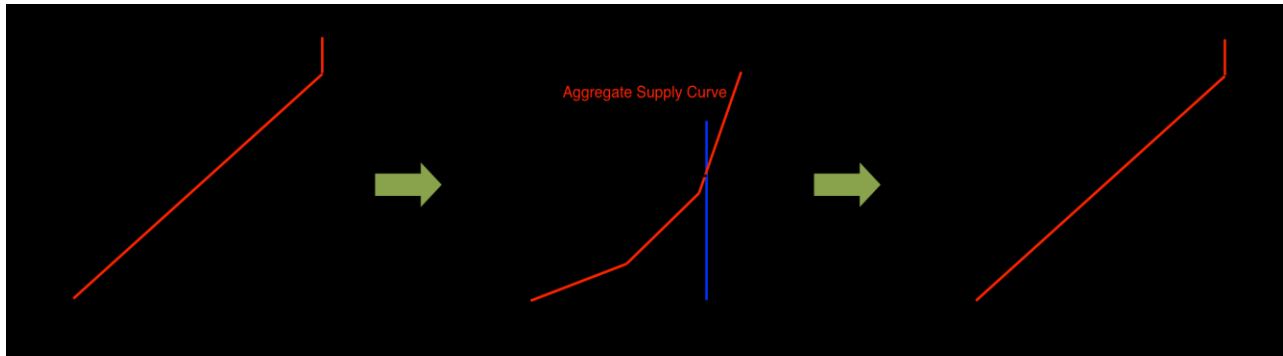
Figure 3.2. Supply curve of a load (left), schematic of market clearing (middle), and determination of quantity of service for the load (right).

To provide a theoretical basis for the above method, we show the connection between the optimality conditions of the optimization problem (2) and the supply curve bidding and market clearing mechanism. The first-order Karush–Kuhn–Tucker optimality condition of the optimization problem requires:

$$C_i'(p_i^*) - \lambda^* - \sum_i(v_{i,1}^* - v_{i,2}^*) = 0, \tag{5}$$

for each i, where $C_i'(p_i^*)$ is the derivative of $C_i(p_i^*)$, λ^* is the optimal Lagrange variable associated with the equality constraint $\sum_i p_i = Q$,, and $v_{i,1}^* \geq 0, v_{i,2}^* \geq 0$ are the Lagrange variables associated with the

inequality constraints $0 \le p_i^k$ and $p_i^k \le \overline{p_i}$, respectively. Moreover, $C_i'(p_i)$ represents the marginal cost of the i^{th} load, which defines its supply curve. We note that λ^* is also the market clearing price.

We next consider a few possible scenarios:

(1) If there is any inequality constraint $-p_i \le 0$ that is active (i.e., $p_i^* = 0, \exists\, i$) and none of the inequality constraint $p_i - \overline{p_i} \le 0$ is active (i.e., $p_i - \overline{p_i} < 0$), we have:

$$C_i'(p_i^*) = \lambda^* + v_{i,1}^*.$$

This implies $\lambda^* = 0$, because $C_i'(0) = 0$, $v_{i,1}^* \ge 0$, and $\lambda^* \ge 0$. Clearly, this means $Q = 0$, and the coordinator does not need any service.

(2) If all of the inequality constraints $p_i - \overline{p_i} \le 0$ are active (i.e., $p_i = \overline{p_i}$ for all i), we have:

$$C_i'(p_i^*) = \lambda^* + v_{i,2}^*.$$

This means $\lambda^* \ge \max_i \overline{\lambda_i}$, and thus each load provides the maximum amount of service. This scenario only happens when $Q \ge \sum_i \overline{p_i}$.

(3) If none of the inequality constraints $-p_i \le 0$ are active, and some of the inequality constraints $p_i - \overline{p_i} \le 0$ are active (i.e., $p_i = \overline{p_i}$ for all $i \in A$, where A is the set of loads whose aforementioned inequality constraints are active), we have
$$C_i'(p_i^*) = \lambda^* + v_{i,2}^*, \quad \text{for all } i \in A,$$
$$C_i'(p_i^*) = \lambda^*, \qquad\quad \text{for all } i \notin A.$$

This means for those loads in set A, the market clearing price is higher than their maximum bidding prices $\overline{\lambda_i}$'s, and hence they provide the maximum amount of service $\overline{p_i}$'s. But still, the sum of their quantities of service is still smaller than the requested amount, i.e., $Q > \sum_i \overline{p_i}$. Therefore, some other loads whose bidding prices are higher must be procured to provide the rest of the amount of service, and the market clearing price is equal to the marginal cost of each of those loads at its optimal power consumption p_i^*. Graphically, this can be explained by Figure 3.2 (middle), which indicates that the optimal price is determined by the intersection of the aggregate supply curve and the vertical line passing the point $(Q, 0)$.

3.3 Simulation-based Demonstration of the Coordination Framework

We next demonstrate the efficacy of the transactive coordination framework through simulations. We consider a supply fan in the AHU, a supply fan in the RTU, and a population of 20 WHs. We assume their cost functions are quadratic,

$$C_i(p_i) = a_i(p_i)^2,$$

where a_i is a positive constant, and p_i is the power increase or reduction compared to the baseline.

We first demonstrate the iterative method based on dual decomposition. We assume the power flexibilities of the supply fan in the AHU, the supply fan in the RTU, and the population of WHs are

$$p_1 \in [-3, 10], \ p_2 \in [-5, 15], \ p_3 \in [-20, 100],$$

and their cost function parameters a_i's are randomly generated between [0.6, 1]. We use the iterative transactive coordination method to coordinate the three types of flexible loads to track the solar generation deviation shown in Figure 2.5. We also solve the optimization problem (2) in a centralized manner to find the optimal solution, and compare it with the solution obtained by the iterative method. Figure 3.3 shows the solar generation deviation, load power deviation obtained from the centralized method, and load power deviation obtained from the distributed iterative method. We see that the optimal solution obtained from the iterative method is the same as the optimal solution that is solved using the centralized method. Moreover, when the solar power deviation is within the power flexibility range of the loads (black dotted lines in Figure 3.3), the load power deviation from baseline can track the solar generation deviation successfully. It is worth mentioning that if the solar power deviation is outside of the flexibility range, the load cannot track it, because it is beyond the flexible loads' response capability. In this case, additional flexible resources are needed. Additionally, we observe that for each time step, about 10 iterations are required for the algorithm to converge using a step size $\alpha = 0.2$. We also find that the convergence rate, which is determined by the step size α, is inversely related to the number of flexible loads. This means that if the coordinator uses the iterative method to coordinate a large number of flexible loads, a large number of iterations will be required for the algorithm to converge at each time step.
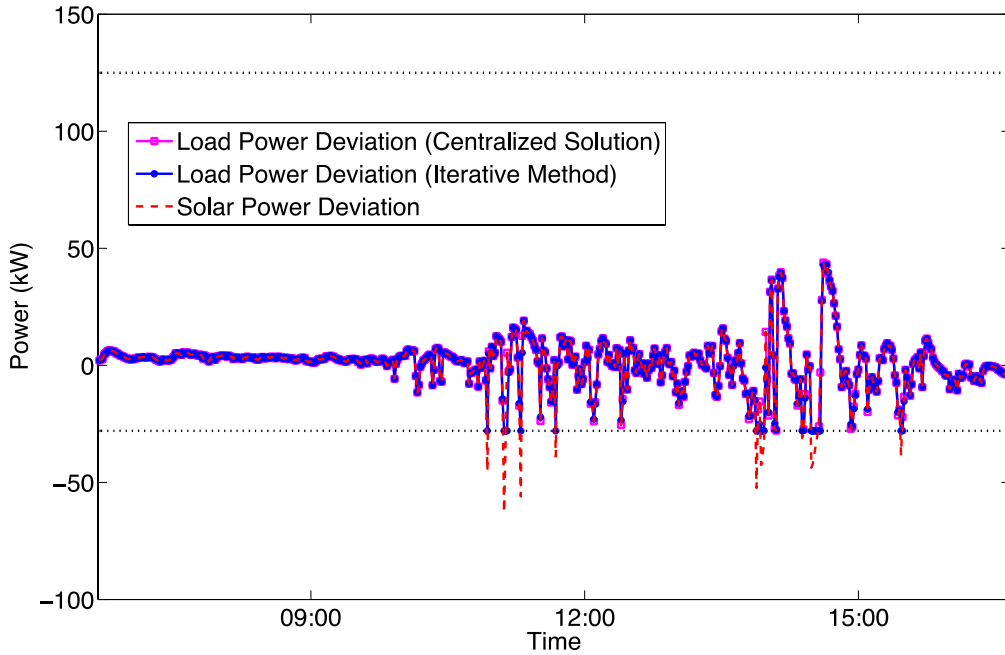
Figure 3.3. Solar generation deviation tracking performance using the iterative method.

We next demonstrate the one-shot supply function bidding method. Because the cost functions are assumed to be quadratic, the marginal cost of each load or the derivative of its cost function is given by

$$C_i'(p_i) = 2a_i p_i.$$

3.5

Graphically, the supply curve submitted by each load is a straight line, and it is depicted in Figure 3.4. We see that the cost function parameter a_i determines the bid price of the provided service. It is worth mentioning that in practice a_i can be chosen as a function of the zone temperature difference from set point, and it might change over time.

At each time step, the supply fan in the AHU, the supply fan in the RTU, and the population of WHs submit a supply curve to the coordinator. Figure 3.4 depicts an example of the three supply curves for the power increase service. The three individual supply curves are then summed into an aggregated supply curve. The intersection of the aggregate supply curve and the vertical line passing the $(Q, 0)$ point determines the market clearing price (MCP). The MCP is then broadcast to the flexible loads, each of which determines its quantity of service based on the received MCP and its own supply curve.
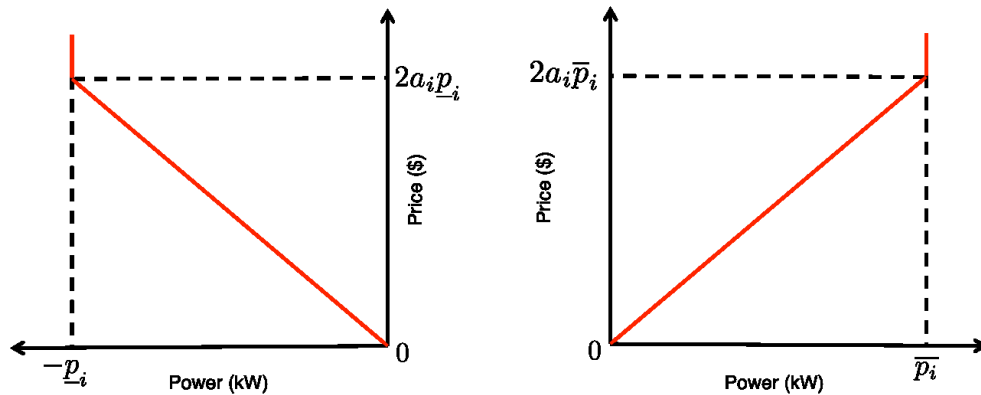
Figure 3.4. Supply curves submitted by the flexile loads.
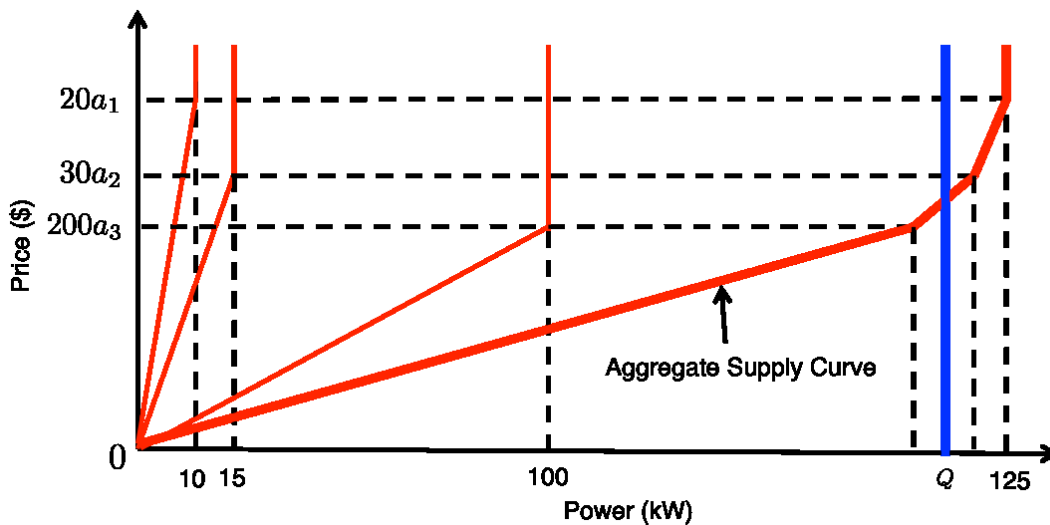
Figure 3.5. Aggregate supply curves and schematic of market clearing.

Figure 3.6 shows the simulation results. In particular, it shows the load power deviations using the supply function bidding method, as well as the centralized solution and the iterative distributed method. We can see that they coincide with each other, which implies that the two different methods yield the same optimal solution at each time step. Moreover, they can successfully track the solar power deviation when it is the within the flexibility range of the loads. Compared to the iterative method, the supply function

3.6

bidding method is able to find the optimal solution in one shot, which significantly reduces the communications between the coordinator and the flexible loads.
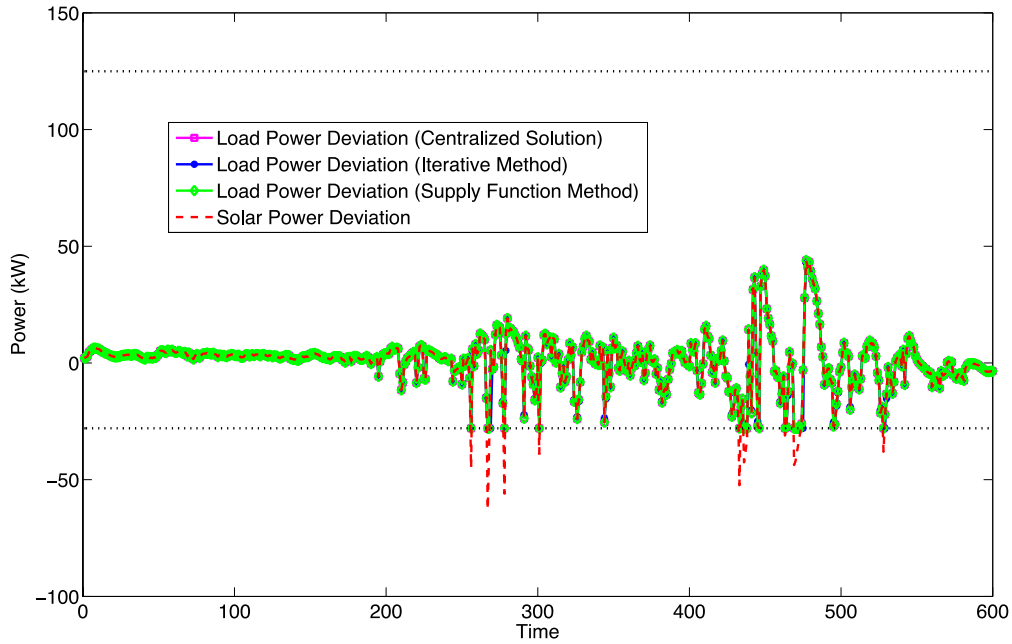
Figure 3.6. Solar generation deviation tracking performance using different methods.

In summary, we found there were in general some tradeoffs between the two different methods. On the one hand, the iterative method based on dual decomposition requires multiple iterations between the loads and coordinator before reaching the optimal solution. However, the only information needing to be exchanged is the price and quantity of the service. On the other hand, the supply function bidding method is able to clear the market in one shot, but each responsive load needs to bid a supply curve, which is a function representing the load's willingness to provide different amounts of services at different prices. The two different approaches are equivalent mathematically. There is a conservation of complexity in the two transactive coordination strategies, in terms of algorithm convergence rate and amount of information needs to be exchanged. However, if the cost functions of the flexible loads are quadratic, the supply curve can be represented by a simple straight line, which can be represented by a point and origin. In this case, the supply function bidding method is more favorable in terms of convergence rate and amount of exchanged information.

The above coordination framework only determines the amount of service each load needs to provide. In the next few sections, we will discuss the device level control strategies that ensure their realized power equals the promised amount of service. In the following sections, we assume the promised amount of service at each time step is given, and it can be treated as a dispatch power profile from the coordinator. The objective of each experiment in the next three sections is to demonstrate control of flexible loads so that their realized power is equal to the power dispatch profile.

4.0   AHU Supply Fans

In this section, we present the results of controlling the supply fan in a commercial building HVAC system to track a dispatched power profile from the coordinator. We first describe the fan model, and identify its model parameters. Next, we present the control strategies, test bed setup, and experiment results. Finally, we summarize the lessons learned from this experiment.
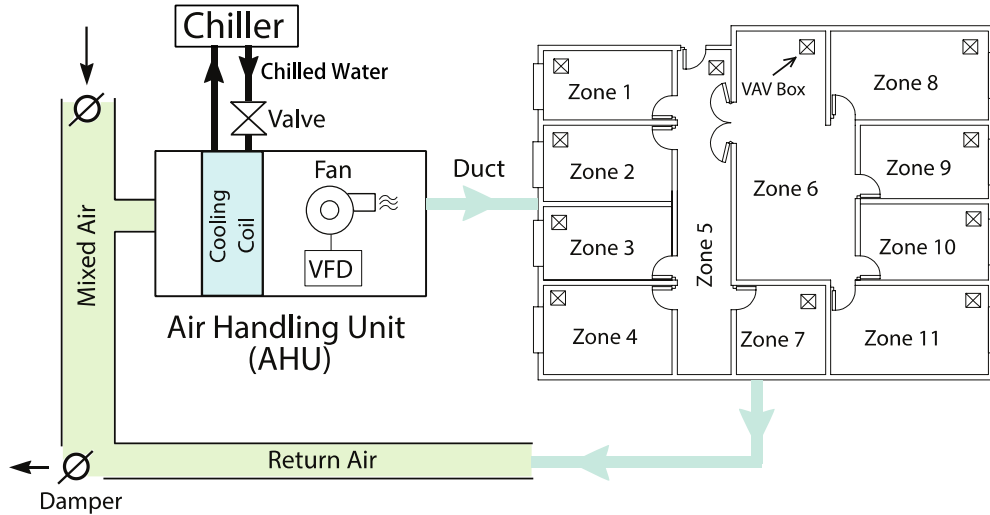
Figure 4.1. Schematic of a multi-zone commercial building HVAC system (Hao, Lin, et al. 2014).

4.1   Fan Model Development

A typical multi-zone commercial building HVAC system is depicted in Figure 4.1. Its main components include the supply fan, cooling coil, variable air volume (VAV) boxes, and chiller. The AHU recirculates the return air from each zone and mixes it with fresh outside air. The ratio of fresh outside air to the return air is controlled by dampers. Mixed air is drawn by the supply fan through the cooling coil in the AHU, which cools the air and reduces its humidity. The air is then distributed to each zone through ducts. The VAV box at each zone controls the airflow rate and supply air temperature going into the zone so that the temperature of the zone tracks a pre-specified desired temperature. As the zonal controllers change the damper position in response to local disturbances (e.g., solar radiation, occupants, etc.), the duct static pressure, which is measured by a sensor that is 2/3 downstream of the duct, changes correspondingly. A fan controller changes the AHU supply fan speed, through a command to the VFD, so as to maintain the static pressure.

The fan power can be modeled as a polynomial function of the total supply airflow (ASHRAE 2008), (Department of Energy 2010),

$$P^f = a_1 v^3 + a_2 v^2 + a_3 v + a_4,$$ (6)

where $\dot{m}$ is the total supply airflow rate of the HVAC system, and a_1, a_2, a_3 and a_4 are constants. Additionally, the supply airflow rate is approximately a linear function of the fan speed,

$$\dot{m} = bv, \tag{7}$$

where v is the supply fan speed, and b is a constant model parameter. Alternatively, the fan power can be modeled as a polynomial function of the fan speed,

$$P^f = c_1 v^3 + c_2 v^2 + c_3 v + c_4, \tag{8}$$

where c_1, c_2, c_3 and c_4 are constant model parameters to be identified.

4.2   System Identification

To develop the fan model and identify its parameters, we use measurement data from one of the AHUs (which serves 17 zones) in a commercial building HVAC system on the PNNL campus. The data are collected by VOLTTRON^{TM} and stored in a Mongo database. We first identify the model parameters a_i's in model (6), using the least squares method with the measurement data of P^f and $\dot{m}$. The identified model parameters are given by $a_1 = 0$, $a_2 = 4.1447 \times 10^{-12}$, $a_3 = -3.1837 \times 10^{-8}$, $a_4 = 2.3303 \times 10^{-4}$, and the R-squared is 0.9602. Figure 4.2 depicts the fan power prediction using the model (2) and the identified parameters, and the fan power measurement. We can see that the fan power prediction matches the fan power measurement very well.
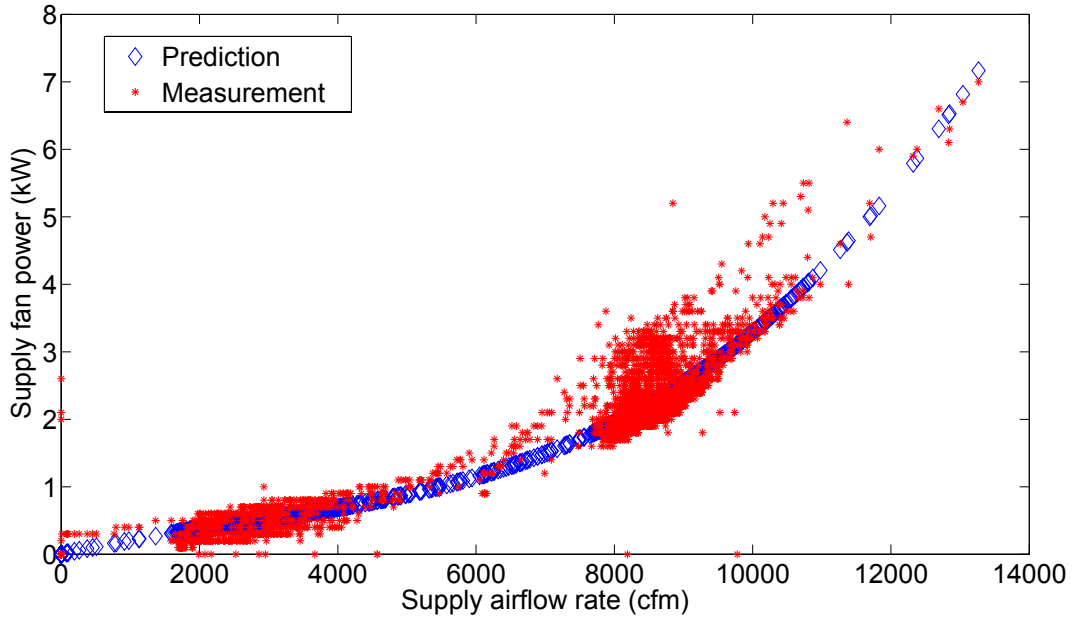
Figure 4.2. Measurement and prediction of fan power as a function of supply airflow rate.

We next identify the model parameter b in model (7). Through the least squares method, we find the parameter b is given by $b = 260.4657$, and the R-squared is 0.9572. The measurement and prediction of supply airflow rate as a function of fan speed is depicted in Figure 4.3. We can see that the supply airflow rate exhibits a linear relationship with the supply fan speed.
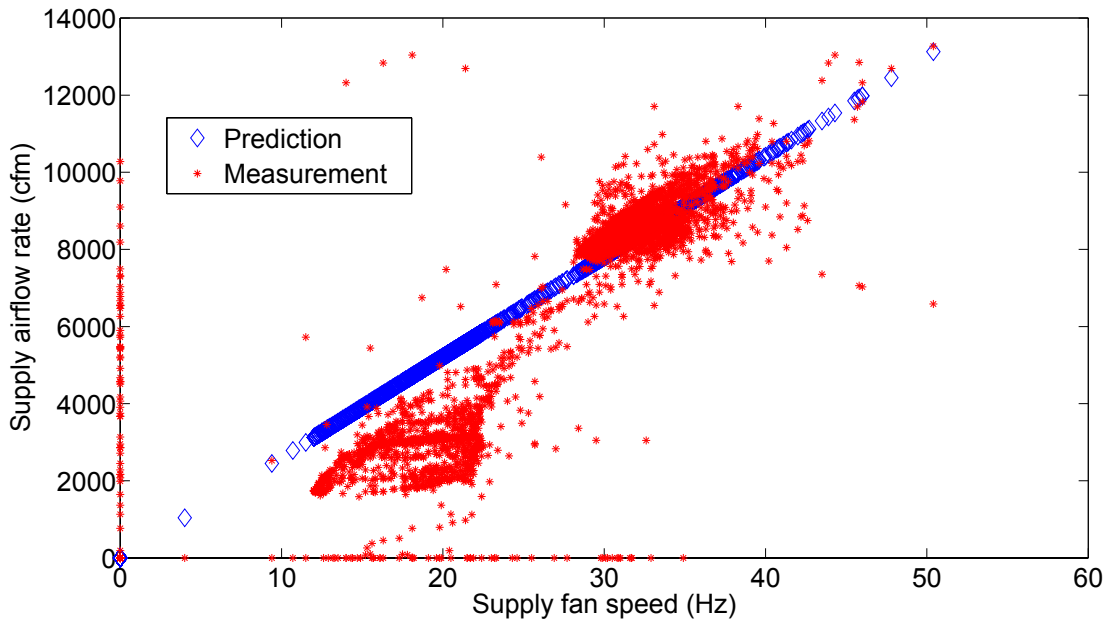
Figure 4.3. Measurement and prediction of supply airflow rate as a function of fan speed.

Finally, we identify the model parameters c_i's in model (8). Similarly, we use the least square method. The model parameters are identified as $c_1 = 0.0044$, $c_2 = 2.8656 \times 10^{-5}$, $c_3 = 0.0022$, $c_4 = -0.0265$, and the R-squared is 0.9835. Figure 4.4 depicts the fan power prediction using model (4) and the identified parameters, and the fan power measurement. We can see that the fan power prediction makes a very good approximation of the fan power measurement. Another interesting aspect we observe from Figures 4.2-4.4 is that the HVAC system operates mainly in two modes, because most of the measurement data are concentrating in two areas.
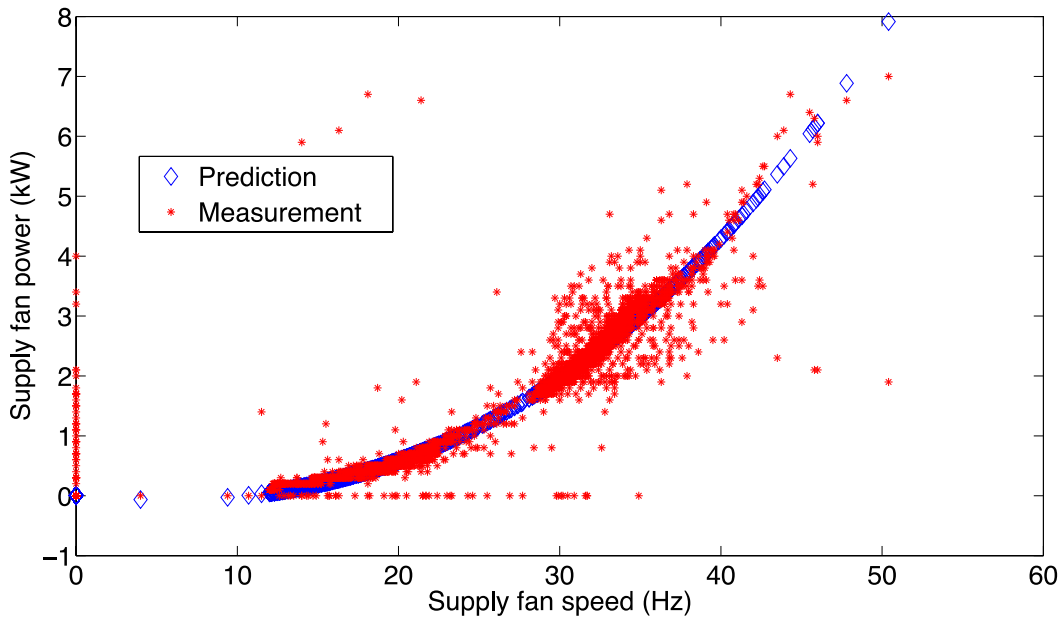
Figure 4.4. Prediction and measurement of fan power as a function of fan speed.

4.3

4.3   Control Strategies

In this section, we present the control strategies we use to control the supply fan to follow a dispatch signal such as renewable generation deviation. In the literature, several control methods have been proposed to control the supply fan power. For example, (Hao, Lin, et al. 2014), (Y. Lin, P. Barooah and S. Meyn, et al. 2015), (Vrettos, Kara, et al., Experimental Demonstration of Frequency Regulation by Commercial Buildings-Part I: Modeling and Hierarchical Control Design 2016), (Vrettos, Kara, et al., Experimental Demonstration of Frequency Regulation by Commercial Buildings-Part II: Results and Performance Evaluation 2016) proposed to control the fan speed and supply airflow rate, and (Maasoumy, et al. 2013) proposed to control the duct static pressure. However, it is our view that the approach taken by (Maasoumy, et al. 2013) is fundamentally flawed, because the fan power is not only a function of duct static pressure, but also a function of supply airflow rate. Moreover, its dependence on supply airflow rate is stronger than the duct static pressure.

In our study, we consider two methods of controlling the fan power as shown in Figure 4.5: 1) control the fan speed (direct controller), and 2) control the duct static pressure (indirect controller). The supply airflow rate is generally not controllable directly in most building automation systems. In the first method, we use model (8) to control the VFD speed directly so that the fan power is commanded to a desired value. It is worth mentioning that if fan power-speed model has a large error, a proportional-integral (PI) controller should be considered to compensate the model-plant mismatch, as proposed by (Vrettos, Kara, et al., Experimental Demonstration of Frequency Regulation by Commercial Buildings-Part I: Modeling and Hierarchical Control Design 2016), (Vrettos, Kara, et al., Experimental Demonstration of Frequency Regulation by Commercial Buildings-Part II: Results and Performance Evaluation 2016). In the second method, in contrast to the proposal in (Maasoumy, et al. 2013), we were not able to develop a fan power-pressure model. Therefore, we consider a model-free feedback PI controller, which measures the current fan power and adjusts the duct static pressure so that the fan power is commanded indirectly to a desired value.
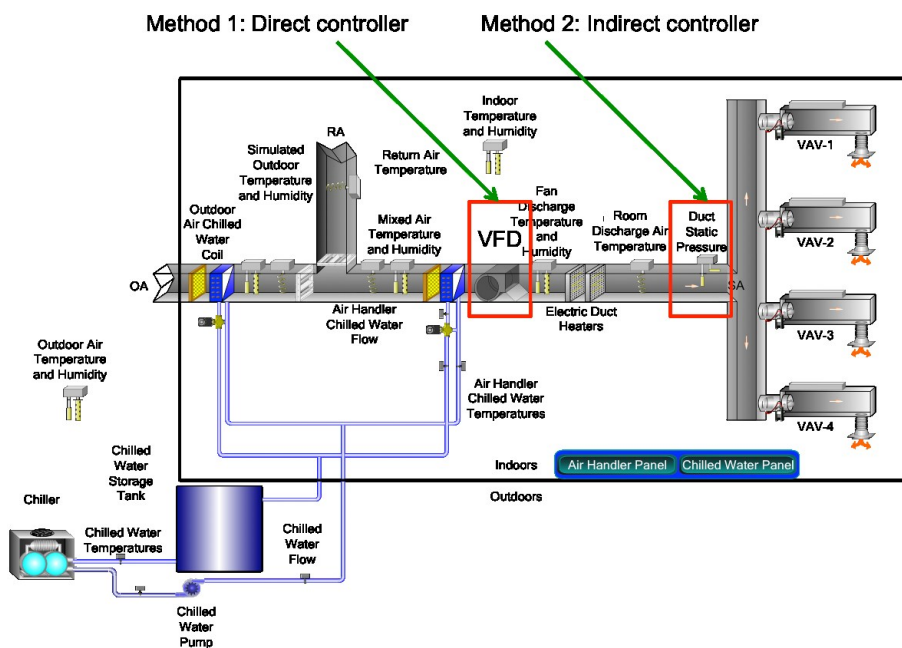

Figure 4.5. The considered two fan controllers in a commercial building HVAC system.

Each method has advantages and disadvantages. The first method has fast response and accurate tracking performance. However, it breaks the existing supply fan control loop (Figure 4.6), in which the fan speed is controlled to maintain a constant duct static pressure. The second method does not need to break the control loop, but it requires adding another feedback control loop to reset the duct static pressure on top of the pressure controller (Figure 4.6). Hence, it has a certain settling time and transient response before the fan power accurately tracks a desired value due to the multiple feedback control loops. Additionally, when tracking a fast changing signal, the indirect controller has less accurate tracking performance than the first method. Both methods work fine for an HVAC system whose baseline fan power is almost a constant or changes slowly during the occupied hour. For an HVAC system whose baseline fan power is difficult to predict, the two methods may have poor performance tracking a dispatch signal. Some examples in the literature use model predictive control to schedule a base operating point for the HVAC fan or chiller (Vrettos, Kara, et al., Experimental Demonstration of Frequency Regulation by Commercial Buildings-Part I: Modeling and Hierarchical Control Design 2016), (Vrettos, Kara, et al., Experimental Demonstration of Frequency Regulation by Commercial Buildings-Part II: Results and Performance Evaluation 2016), (Mai and Chung 2015), (Hao, Lian, et al. 2015), (Lin, Barooah and Mathieu, Ancillary services to the grid from commercial buildings through demand scheduling and control 2015). However, these methods need to vastly change the existing building control loop and require accurate zone thermal models and load models, which are not desirable for building operators. Moreover, sustainable maintenance of these MPC controllers is a challenge for the building operators. Nevertheless, for general demand response purposes, whose goal is to reduce the power consumption of an HVAC system for a period of time, the two methods are good enough to curtail the load of commercial buildings.
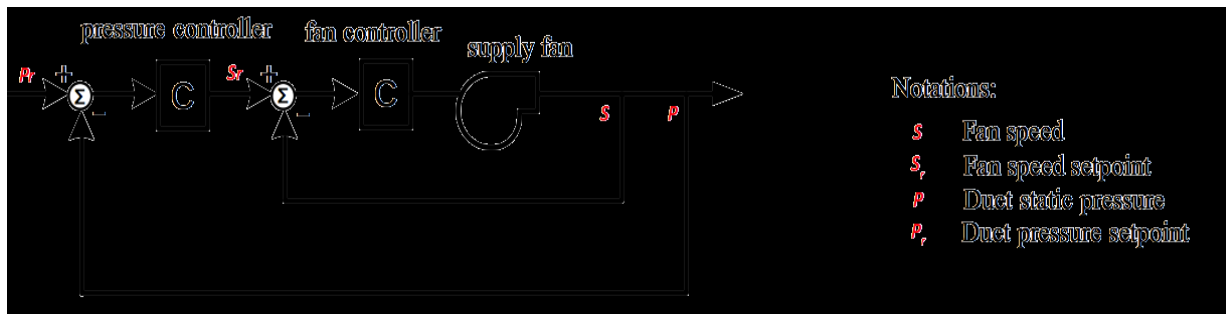
Figure 4.6. Schematic of existing supply fan control in a commercial HVAC system.

4.4 Test Bed Setup

In this section, we present the test bed used in the experiment demonstration. We consider a 4-zone commercial building HVAC system with 1 VFD-controlled supply fan and 4 VAV boxes (which serve the same room). The system schematic is depicted Figure 4.5. Figure 4.7 shows a few pictures of the test bed. The test bed setup is based on VOLTTRON™—a low cost and open-source software platform (Pacific Northwest National Laboratory 2016). The VOLTTRON™ platform provides basic agents (e.g., listener agent, actuator agent, weather agent, etc.), which facilitate commination with and control physical HVAC equipment and developing new agents to enable customized applications. The VOLTTRON-based test platform is depicted in Figure 4.8. In our experiments, we used the following six agents:

- Proxy Agent. The BACnet proxy agent provides communication with BACnet device in an HVAC system. If the HVAC system is using Modbus protocol, then the BACnet proxy agent is not needed.

- Master Driver Agent. The master driver agent configures the device, and it is the key for communicating with devices using VOLTTRON™.
- Listener Agent. The listener agent "listens" to all topics published on the common message bus. It is useful to testing if agents under development are publishing message correctly.

Figure 4.7. Pictures of the HVAC system used for the experiment.

- Historian Agent. The historian agent is an agent storing the messages published on the message bus in a database. Currently, SQLight, MySQL, and sMap databases are supported.
- Actuator Agent. The actuator agent can access the control points of a physical device controller. It may request scheduled times to interact with one or more physical devices.
- Control agent. The two controllers proposed in this report are programmed using Python as control agents to control the supply fan in the HVAC system.

The VOLTTRON™ platform has many great features. For example, it is open, flexible, and modular. All of its major functionalities are programmed as agents with basic properties and functions, which can be reused later. Additionally, it is a very low-cost platform that is easy to deploy. The only hardware requirement is a single-board computer such as Raspberry PI or Beaglebone. Moreover, it supports remote control. As long as the appropriate Internet Protocol (IP) address and authorization are given, the VOLTTRON™ platform is able to connect with and control the physical HVAC device from anywhere using the Internet.
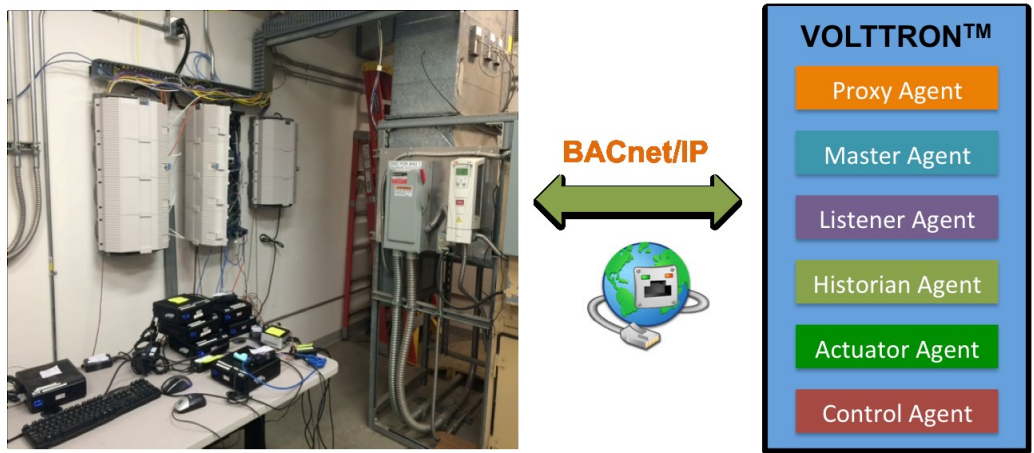
Figure 4.8. Configuration of the VOLTTRON™ -based test bed.

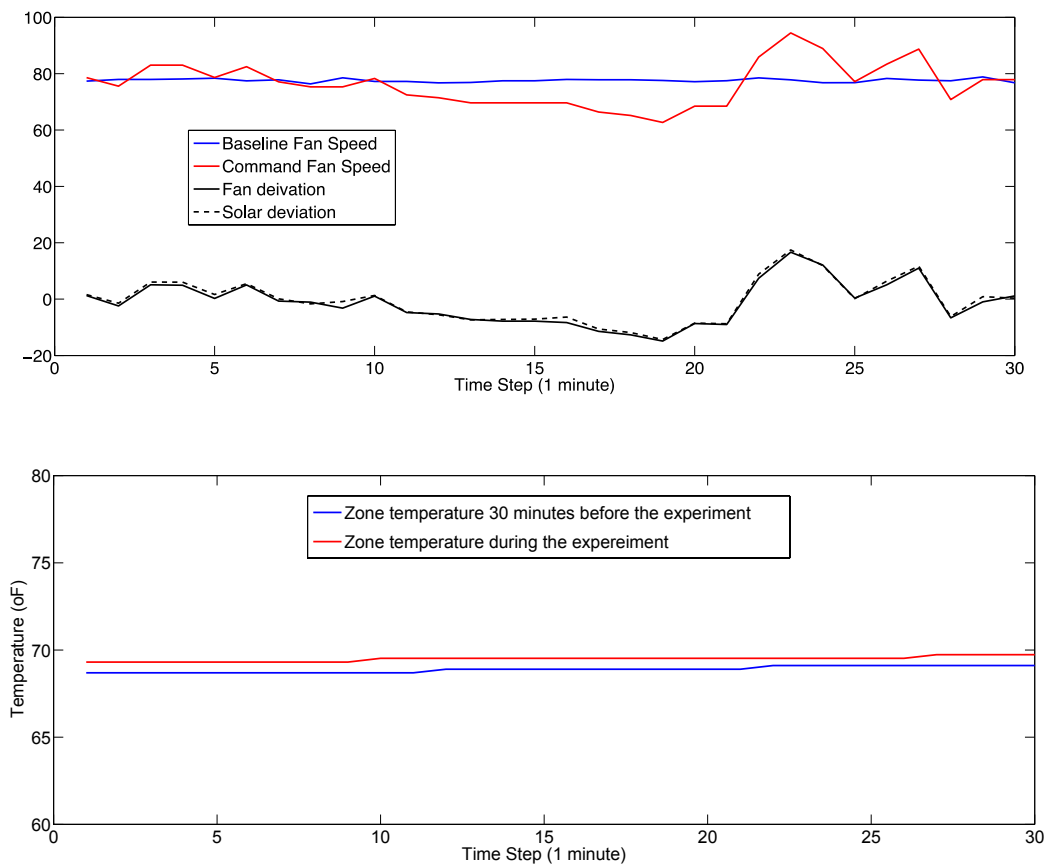## 4.5   Experiment results

Figure 4.9. Experiment result of using the direct controller.

In this section, we present experimental results using both the direct and indirect controllers we proposed. Because the HVAC system used for this experiment does not have accurate power measurement, we therefore control the fan's speed to track a desired speed trajectory, rather than a power trajectory. Figure

4.9 shows the experiment results using the direct controller. The top figure shows the baseline fan speed, realized fan speed, fan speed deviation from baseline, and solar deviation. We can see that the fan speed deviation can quickly track the solar deviation very well. Moreover, the zone temperatures 30 minutes before the experiment and during the 30-minute experiment are shown in the bottom figure. We observe that the zone temperature was increasing very slowly. However, it is our view that zone temperature increase is not due to our experiment, because the zone temperature is rising slightly before initiation of the experiment. This is probably because the zone temperature is floating in a deadband, and the outside temperature is increasing.

We next present the experiment results using the indirect controller, which controls the duct static pressure to indirectly command the fan speed to four different values (98%, 88%, 78%, and 95%). First of all, we use 1-minute sampling time for feedback. Figure 4.10 shows the static pressure, static pressure set point, fan speed, and zone temperature during the experiment. We can see from the third subfigure that the fan speed is able to track the four different values, but with slow transient response, which is not very desirable.
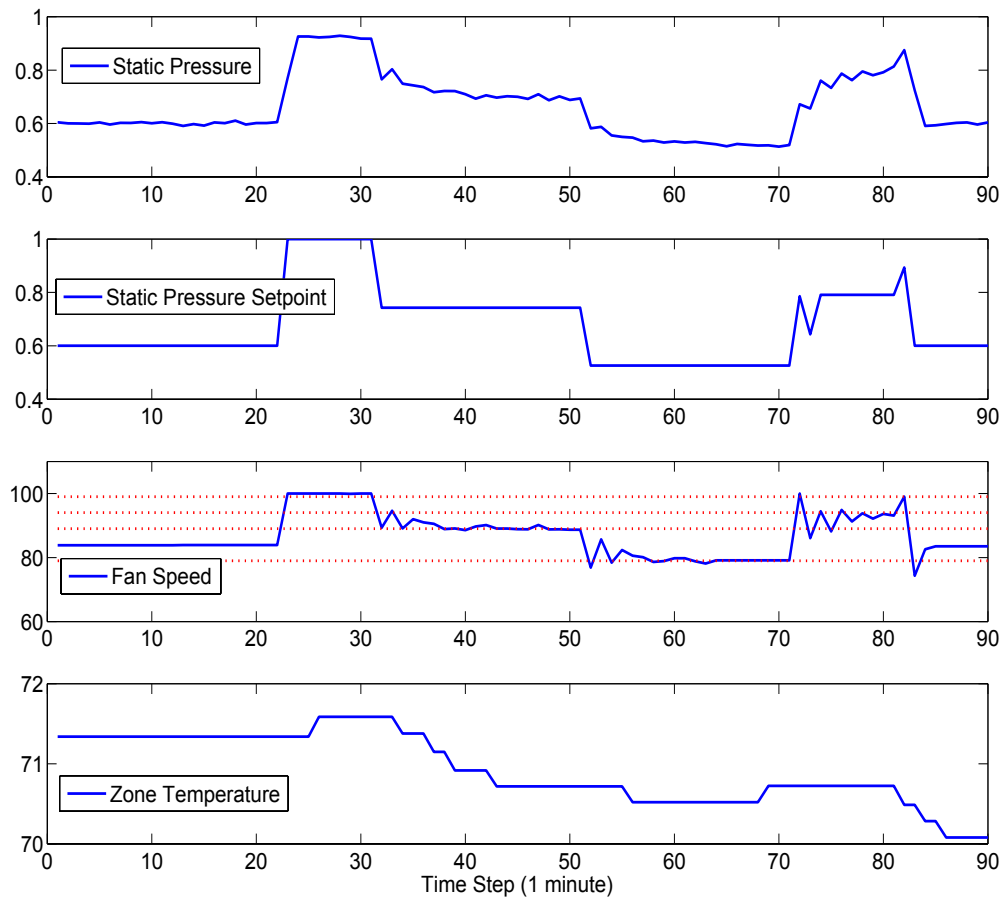
Figure 4.10. Experiment results of using the indirect controller (time resolution 1 minute).

To have fast and accurate tracking performance, we increase the sampling time for feedback from every 1 minute to every 10 seconds. Figure 4.11 illustrates the experiment results with the higher sampling rate. We can see from the third subfigure that the tracking performance is improved significantly. Moreover,

the temperatures in both Figure 4.10 and Figure 4.11 do not change much during the experiments. This implies that the experiments have no significant impact on the zone thermal comfort. Using different sampling rates requires different PI gains, and tuning the PI gains appropriately will yield more accurate tracking performance. In our experiment, we use Ziegler-Nicholes method (Wikipedia 2016) to find the initial PI gain values, and use the trial and error method to refine the control gains.
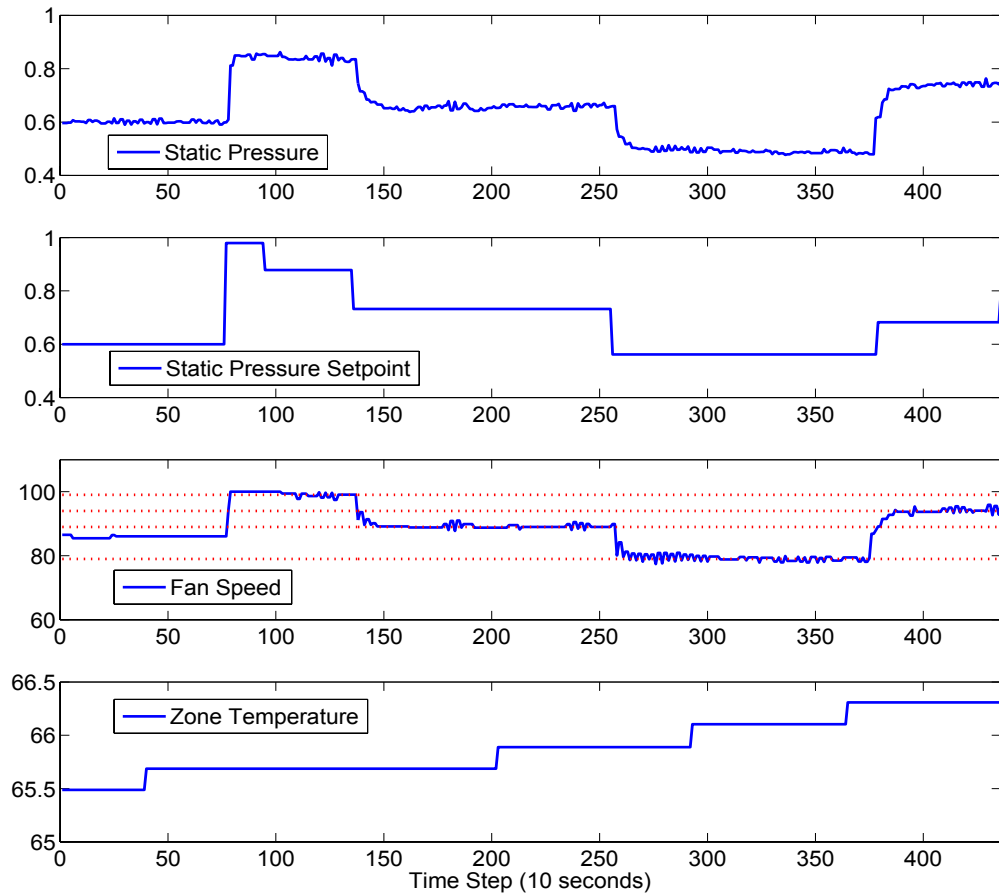

Figure 4.11. Experiment result of using the indirect controller (time resolution 10 seconds).

In summary, we show that our control methods are able to command the supply fan to track a desired trajectory, while having minor impact on the zone thermal comfort. In the future, we will conduct additional experiments on other commercial HVAC systems, and present the findings and lessons learned.

5.0 Packaged Rooftop Units

In this section, we study modeling and control of a simulated packaged RTU to track a dispatch power command. We present the system description and model identification. We also discuss its operating mode and legacy RTU control logic. An RTU simulation test bed is set up using Modelica and Dymola, and a couple of case studies are presented. We show that RTU had a great potential for tracking a power dispatch signal from the grid, while having little impact on the thermal comfort of the serving zone.

5.1 System Description and Identification
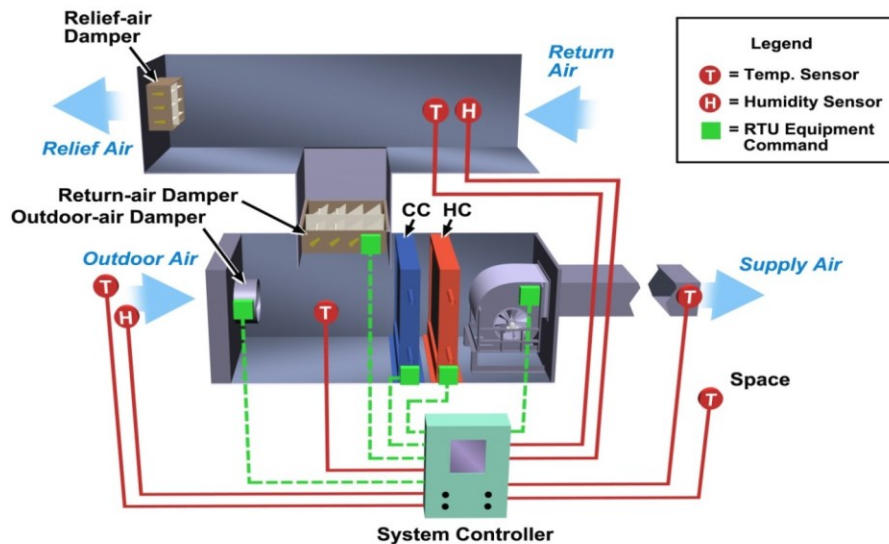
Figure 5.1. Schematic diagram of a packaged rooftop unit.

Packaged RTUs are used in about 50% of all commercial buildings, and they serve over 60% of the commercial building floor space in the United States (Energy Information Administration 2016), (Wang, et al. 2013). The configuration of a typical RTU is depicted in Figure 5.1. Its main components include the supply fan, heating coil, cooling coil, and outdoor-air, return-air, and relief-air dampers. The movement of the airflow is described as follows. The indoor air from the zones is circulated to the HVAC system through ducts. Some of this flow (relief air) is exhausted outside of the building through relief-air dampers, while the other portion (return air) is directed towards the RTU via the return-air damper. The intake of the outdoor air is controlled by the outdoor-air damper. The outdoor air is then mixed with the return air and conditioned by the cooling/heating coil to appropriate temperature and humidity set points. Afterward, the supply air is drawn by the supply fan, and provided to the zones through ducts.

In our experiment, we aim to control the supply fan speed so that its power can track a value dispatched by the coordinator. The supply fan only runs when the building is occupied, and its fan speed is generally varied based on the operation modes of the RTU, which are described as follows:

(1) Ventilation mode
- When there is no call for heating or cooling, the system will be in the ventilation mode and the fan will run at 40% of its full speed (adjustable).

(2) Cooling mode
- When the outdoor dry bulb temperature is less than 58ºF or greater than 70ºF
  - If there is a call for first-stage cooling, the fan will run at 75% of its full speed (adjustable).
  - If there is a call for second-stage cooling, the fan will ran at 90% of its full speed (adjustable).
- When the outdoor dry bulb temperature is greater than 58ºF but less than 70ºF, the fan will ran at 90% of its full speed (adjustable) if there is a call for either first or second-stage cooling.

(3) Heating mode
- When there is a call for first-stage heating, the fan will run at 75% of its full speed (adjustable).
- When there is a call for second-stage heating, the fan will run at 90% of its full speed (adjustable).
- When there is a call for auxiliary electric heat on a heat pump, the fan will run at 90% of its full speed (adjustable).

In addition to the above modes, the fan speed will be overridden under the following conditions:

- In the cooling mode, if supply air temperature drops below the low limit set point of 45°F, the fan will run at 100% of its full speed.

- In the heating mode, if the supply air temperature rises above the high limit set point of 135°F, the fan will also run at 100% of its full speed.

- If the carbon dioxide (CO_2) concentration is controlled to keep the space CO_2 levels of below the CO_2 set point (1000 PPM default) by a PI loop, the fan speed will modulate after the damper is at 100% open.
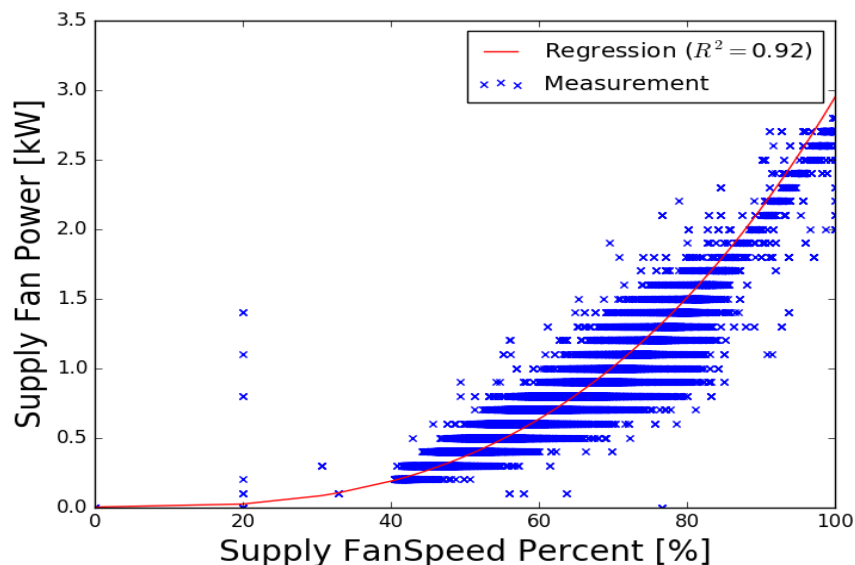

Figure 5.2. Prediction and measurement of fan power as a function of fan speed ratio.

Again, we model the fan power as a polynomial function of the fan speed,

$$P^f = c_1 v^3 + c_2 v^2 + c_3 v + c_4,$$ (9)

5.2

where c_1, c_2, c_3 and c_4 are constant model parameters to be identified. We use the least square method to identify the parameters, and get $c_1 = 2.9573 \times 10^{-6}, c_2 = -4.5 \times 10^{-7}, c_3 = -8.1 \times 10^{-5}, c_4 = 3.0 \times 10^{-4}$, and $R^2 = 0.92$ based on the measured data of a supply fan that serving a PNNL building. Figure 5.2 depicts the fan power prediction using model (9) and the identified parameters, and the fan power measurement. We can see that the fan power prediction matches the fan power measurement fairly well. One of the great advantages of controlling RTUs for grid service is that the baseline fan speed is determined based on the operational modes of the RTU. Once we know the operational mode, we can use the VFD to modulate the fan speed based on model (9), so that its power consumption minus baseline tracks a dispatched power command. Additionally, if there is a large model-plant mismatch, a closed-loop PI controller can be used to adjust the fan speed more accurately to improve the tracking performance.

5.2   Test Bed Setup

In our experiment, we consider one RTU that is modeled with Modelica (Modelica Association 2014). Modelica is an object-oriented modeling language. The developed RTU model is compiled and simulated in Dymola (Dassault Systemes 2016) (which is the simulation environment for Modelica).
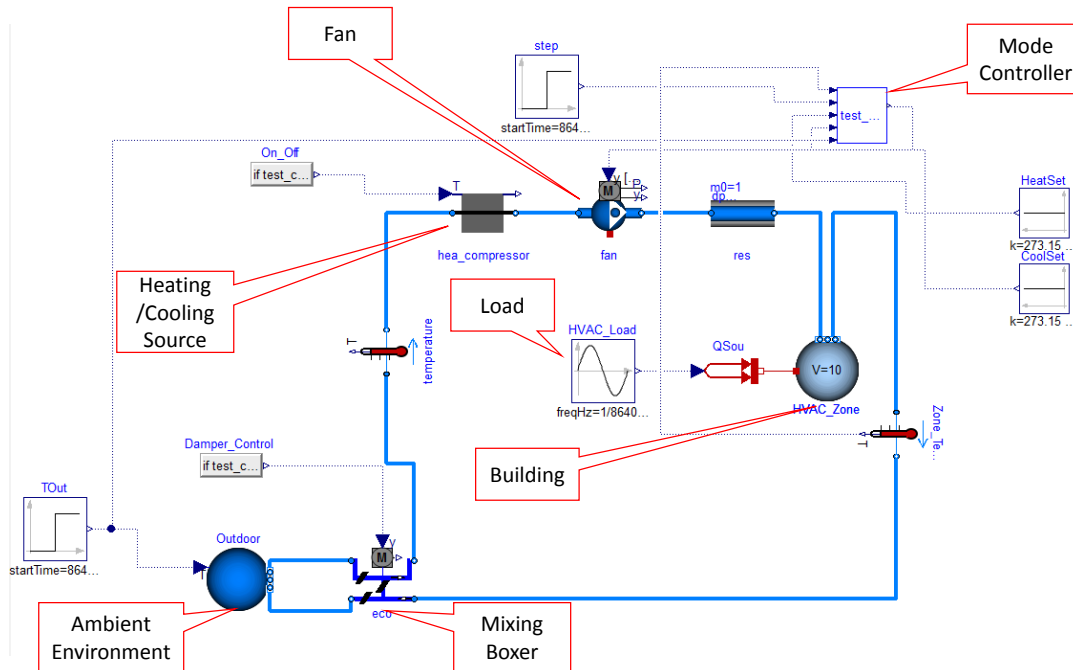
Figure 5.3. The diagram of the top-level RTU model.

Modelica enables us to develop a hierarchical model for the RTU. Figure 5.3 shows the diagram of the top-level RTU model, which consists of sub-models for physical components of an RTU (such as the heating/cooling source, fan, mixing box), a mode controller, and an HVAC zone. The sub-models for the physical components of the RTU and the HVAC zone are built based on the component models from the Modelica *buildings* library 3.0.1 (Wetter, et al. 2015) (see Table 5.1 for details). The measured data mentioned in Section 5.1 are used to calibrate the fan model. The volume for the HVAC zone is set as 1000 m³. The sub-model for the mode controller is built with *Modelica_StateGraph2* library 2.0.3 (Otter, et al. 2009). Figure 5.4 is a diagram of the sub-model for the mode controller. This sub-model takes

account of six operating modes of the RTU: the first-stage mechanical cooling, the second-stage mechanical cooling, the ventilation, the economizer, the first-stage heating, and the second-stage heating.

Table 5.1. The component models from the Modelica Buildings library to build the simulation test bed.

Components	Models from Modelica Buildings library
Fan	Fluid.Movers.SpeedControlled_y
Heating/Cooling Source	Fluid.HeatExchangers.HeaterCooler_
Mixer Boxer	Fluid.Actuators.Dampers.MixingBox
HVAC zone	Fluid.MixingVolumes.MixingVolume
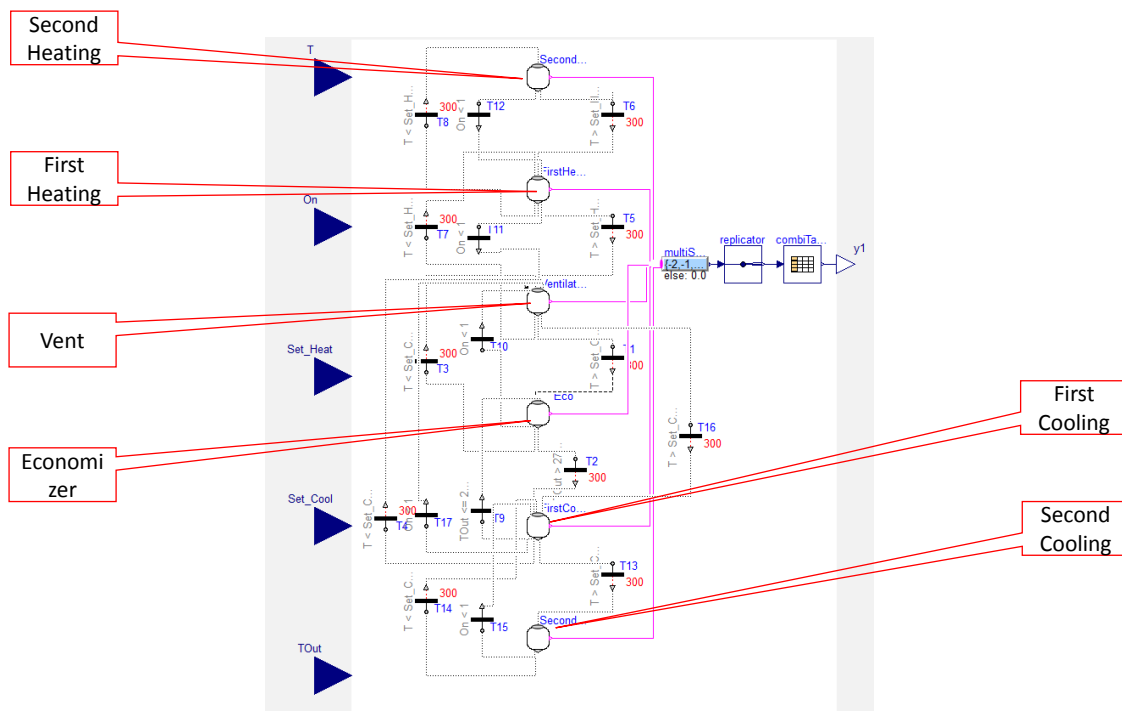
Figure 5.4. Diagram of the mode controller model.

5.3   Experiment results

In our simulation study, we consider two scenarios. Both scenarios cover a period of 86,400 seconds (one day). In scenario 1, the RTU is always in the second-cooling mode and the cooling load is set as a constant of 22 kW. As shown in Figure 5.5 (top), the fan power deviation tracks the dispatch power command quite well. Figure 5.5 (bottom) shows the trajectories of the zone temperature with and without tracking of the dispatch power command. We can see that the baseline zone temperature is almost constant during the experiment, because its cooling load is assumed to be a constant. When tracking the dispatch power command, we observe that there are some oscillations in the zone temperature. However, the amplitudes of the oscillations are less than 0.5°C. This means tracking the dispatch power command does not significantly affect the thermal comfort of the studied HVAC zone.

In scenario 2, we consider a more realistic case, in which we use a sinusoid signal as a synthetic cooling load for the studied zone. The frequency, amplitude, and offset of the sinusoid signal are $5.8×10^{-6}$, 3 kW,

and 18 kW, respectively. Under such a cooling load, the operating mode of the RTU switches between the first-stage cooling and the second-stage cooling, as shown in Figure 5.6. Additionally, we can see that the fan power deviation tracks the dispatch power command very well and the zone temperature has no significant deviations from the baseline.

In the future, we are interested in implementing the designed control strategies using VOLTTRON on a physical RTU test bed, and examining the potential of tracking a dispatch power signal and its impact on RTU operation and thermal comfort to the occupants.
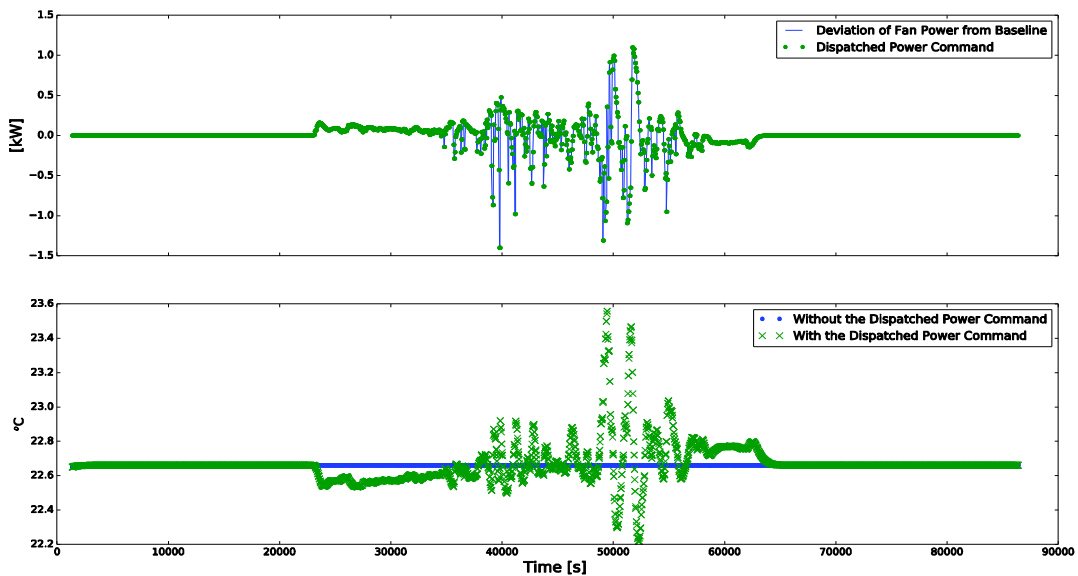

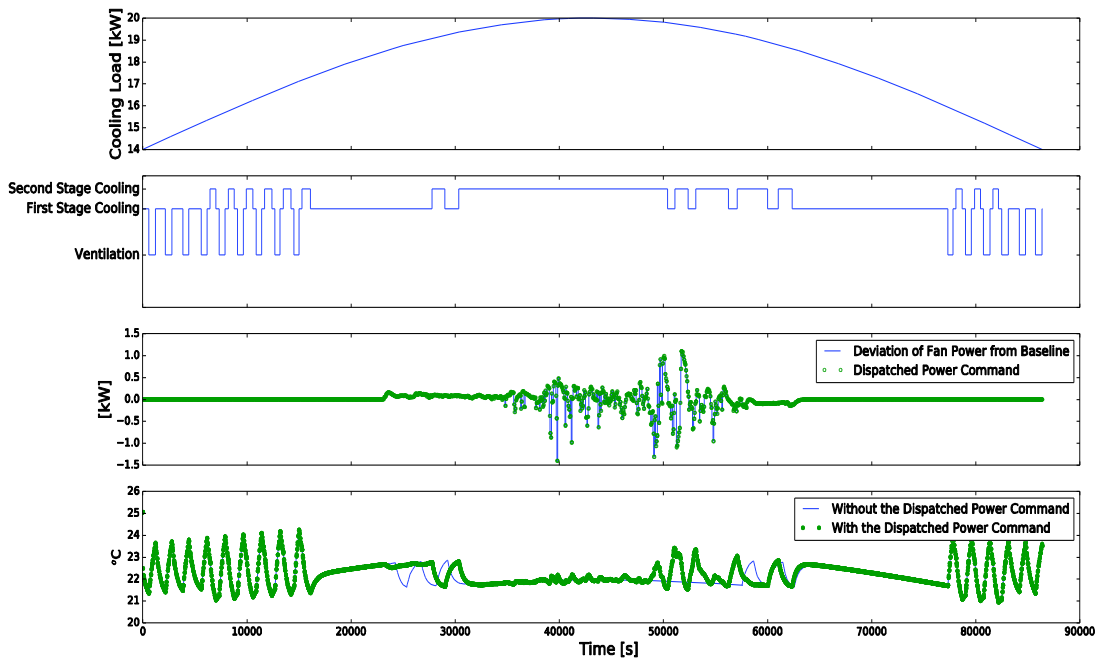Figure 5.5. Simulation results for scenario 1: (top) fan power and (bottom) zone temperature.

Figure 5.6. Simulation results for scenario 2: (from top to bottom) cooling load, RTU operating mode, fan power, and zone temperature.

6.0 Water Heaters

In this section, we study coordination of a population of WHs for renewable integration. We develop a simple first-order WH model and identify the model parameters. We describe the control system architecture and present a priority-stack control algorithm to turn WHs ON/OFF. We also present the test bed setup and discuss the experiment results. Using a HIL simulation, we demonstrate that coordinating WHs is very effective in absorbing the uncertainty of renewables while having small impacts on the end-users.

6.1 Problem Statement

TCLs such as WHs, air-conditioners, heat pumps, and refrigerators, have a great potential for renewable integration because of their huge inherent thermal storage capacity. In this section, we present a HIL demonstration of aggregation and control of a large population of domestic WHs to track solar generation deviation. In practice, each WH's power consumption is discrete; it is either the rated power or zero. To get a smooth response, we will consider a large population of them to deliver a smooth aggregate response. However, in practice, it is very challenging to engage a large number of WHs for experimental purposes. Therefore, we consider a combined hardware and model in the loop demonstration, in which most of the WHs will be simulated using models in MATLAB, while having a physical WH. A schematic of the HIL demonstration is provided in Figure 6.1. The WH models and control algorithms reside in MATALB. The VOLTTRON^{TM} platform collects data from the solar PV panel and the physical WH, and communicates with MATLAB. The VOLTTRON^{TM} platform sends control commands to the physical WH to turn it ON or OFF.

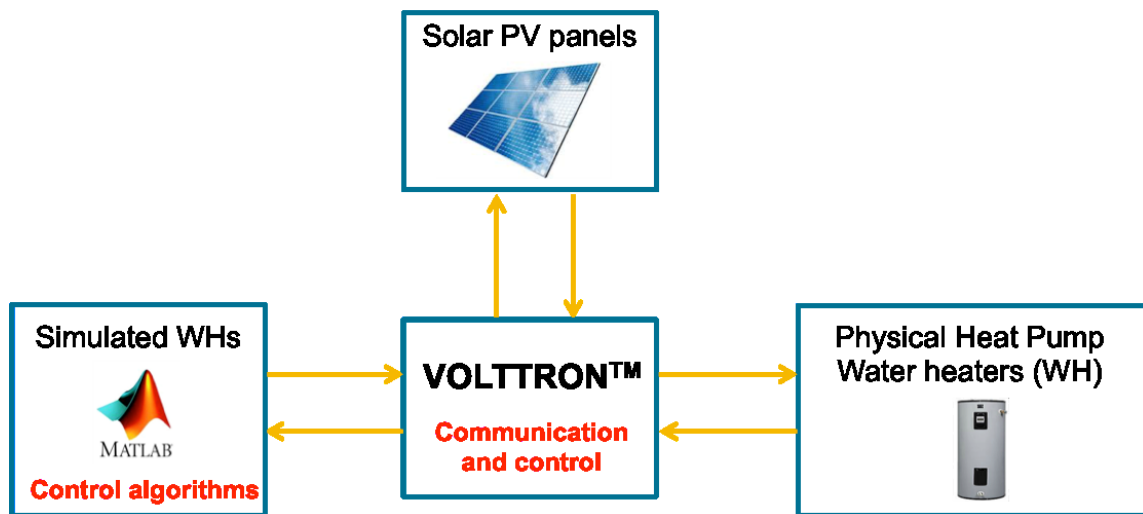
Figure 6.1. Schematic of HIL demonstration of controlling WHs for renewable integration.

6.2 Model Development

The temperature of each WH can be described by a hybrid system model:

$$C\frac{d\theta(t)}{dt} = \frac{\theta_a - \theta(t)}{R} + \begin{cases} P + w(t), & \text{ON State} \\ w(t), & \text{OFF state} \end{cases},$$ (10)

where $\theta(t)$ is the WH tank temperature, θ_a is the ambient temperature, C and R are its thermal capacitance and resistance respectively, P is its rated power when it is ON, and $w(t)$ represents water draw by the end user. Each WH has a temperature set point θ_r, and it cycles ON and OFF between a temperature band $[\theta_r - \Delta, \theta_r + \Delta]$ (Figure 6.2),

$$\text{status} = \begin{cases} \text{ON}, & \theta(t) < \theta_r - \Delta, \\ \text{OFF}, & \theta(t) > \theta_r + \Delta. \end{cases}$$

6.3 System Identification

We use measured data collected from a WH on the PNNL campus. The least squares method is used to identify the model parameters, which yields $C = 0.0838$ kWh/°C, $R = 494.3283$ °C/kW, and $P = 4.0700$ kW. Figure 6.2 depicts the prediction and measurement of WH temperature. We see that the temperature prediction matches very well with the temperature measurement. The identified model parameters are used to simulate a population of WHs in MATLAB for demonstration purposes.
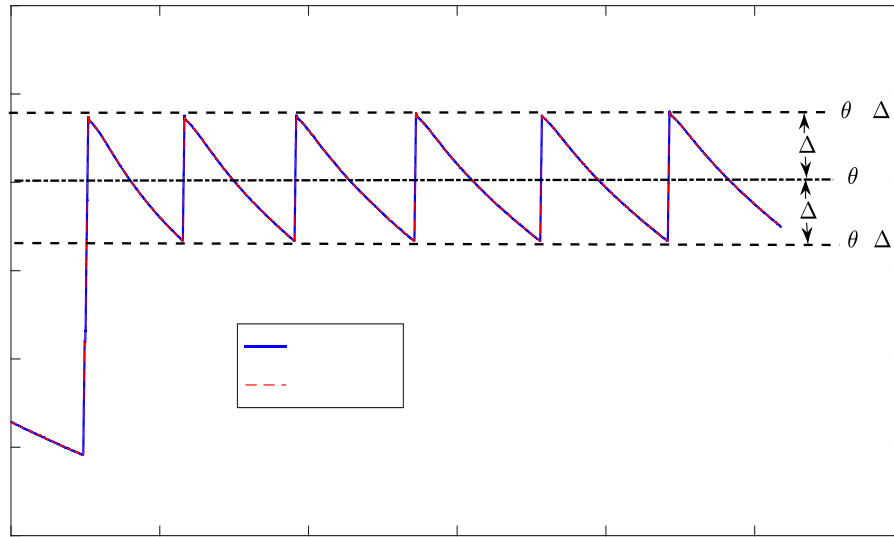
Figure 6.2. Comparison of prediction and measurement of WH temperature.

6.4 Control Strategies

In our study, we consider a centralized control scheme to aggregate and control WHs. The schematic of the control system is shown in Figure 6.3. At each time step, we predict the aggregate power consumption of the population at the next time step, and compare it with the baseline power of WHs. If the difference between the aggregate power and baseline power is larger than the predicted renewable generation

deviation at the next step, then we turn off some of the WHs. Similarly, if the difference is smaller than the renewable generation deviation, we turn on some of the WHs. To minimize the ON-OFF switching times for each WH, we use a priority-stack controller (Hao, Sanandaji, et al. 2015), (Hao, Sanandaji, et al. 2013) to decide which WHs to turn on or turn off.
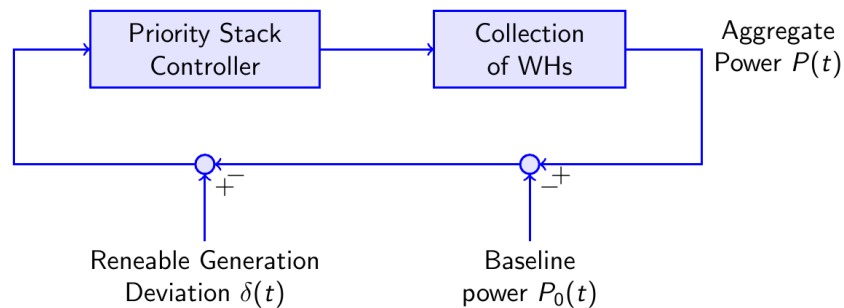
Figure 6.3. Schematic of overall control structure.

The priority-stack controller is depicted in Figure 6.4. For each state of WHs, we build a priority stack, where the priority is measured by the difference between its current temperature and the appropriate temperature bound. For example, for WHs that are OFF, the priority is calculated by the distance between its current temperature and its lower temperature bound $\underline{\theta} = \theta_r - \Delta$. The smaller the distance, the higher priority it receives to be turned on. In this way, we make sure we turn on the WH that has been turned off for the longest time first and turn off the WH that has been turned on for the longest time first, to prevent frequent ON-OFF switching to protect the hardware. More details about the priority-stack-based controller can be found in (Hao, Sanandaji, et al. 2015) and (Hao, Sanandaji, et al. 2013).

Figure 6.4. Priority-stack controller.

6.5   Test Bed Setup

In our experiment, we consider a population of 1000 WHs, where 999 WHs are simulated in MATLAB using the model parameters identified in Section 5.3 and 1 of the physical WHs on the PNNL campus, as

shown in Figure 6.5. The control system and priority-stack controller are implemented in MATLAB. The HIL demonstration is executed as follows: at each time step, MATLAB predicts 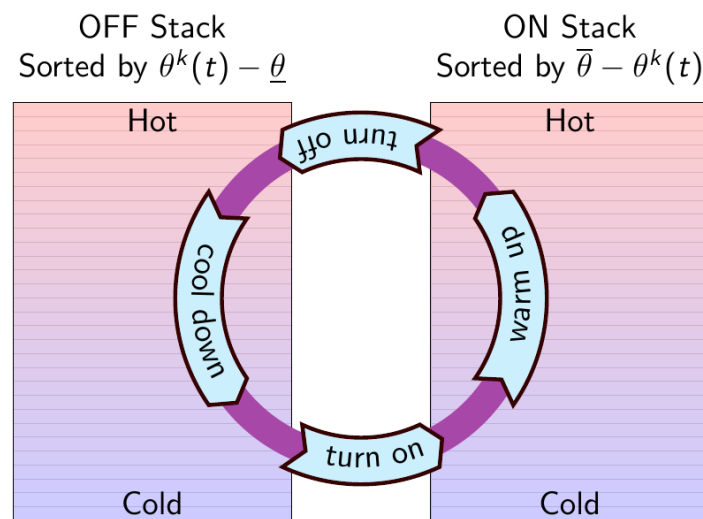the aggregate power of the population of WHs using the developed models in Section 5.2 based on the current WH temperature and ON/OFF status. It then calculates the power difference from the baseline of the WHs. At the same time, MATLAB receives solar generation data, and predicts the solar power output at the next step. On the basis of the power difference and the solar generation deviation, the priority-stack controller determines which WHs are to be turned ON or OFF. The control commands are then sent to the simulated WHs in MATLAB, and the physical WH via VOLTTRON™. The experiments repeats this process until the end of the time horizon considered.

Figure 6.5. Picture of physical WHs.

6.6 Experiment results

We conduct the experiment for 24 hours. We consider two cases: 1) the solar power prediction is obtained by taking the 15-minute trailing average of the historic solar generation data; and 2) the solar power prediction is taken as the ideal solar production on a clear-sky day. Considering the first case, the solar power deviation and WHs power deviation from baseline are depicted in Figure 6.6 (a). We see that the WHs power deviation from baseline can successfully track the solar generation deviation. Moreover, Figure 6.6 (b) and (c) show that the temperatures of simulated and physical WHs can be well regulated within the user-specified temperature bounds. This means tracking the solar generation deviation has little impact on the WH temperature. Furthermore, we examine the impact of our controller on the ON-OFF switching times of WHs, and their total energy consumption. We find the average switching time of all WHs during the 24-hour period with and without our control are 6.4220 and 6.1840, respectively. Moreover, their total daily energy consumptions with and without our control are 1856 kWh and 1835 kWh. This means tracking the solar generation deviation (in case 1) has little impact on the quality of service to the end-users.
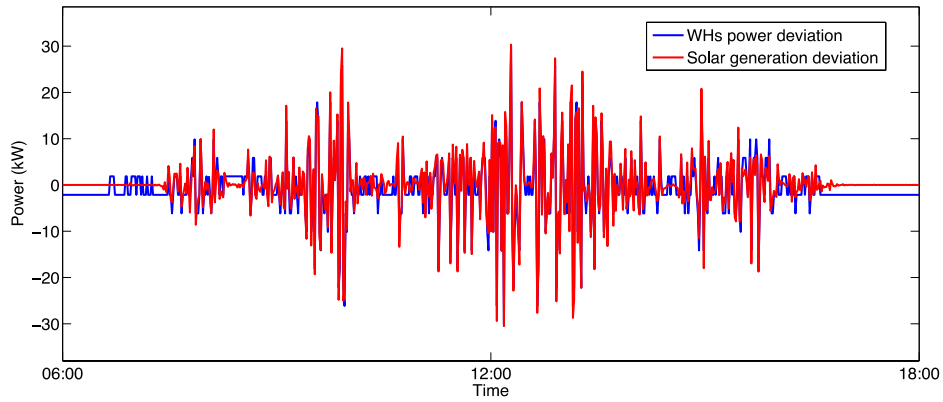
(a) Power-tracking performance
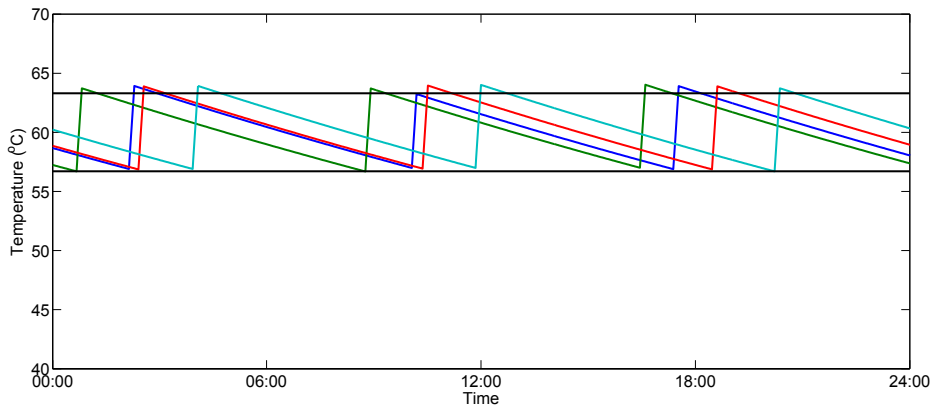
(b) Temperature trajectories of a few simulated WHs.
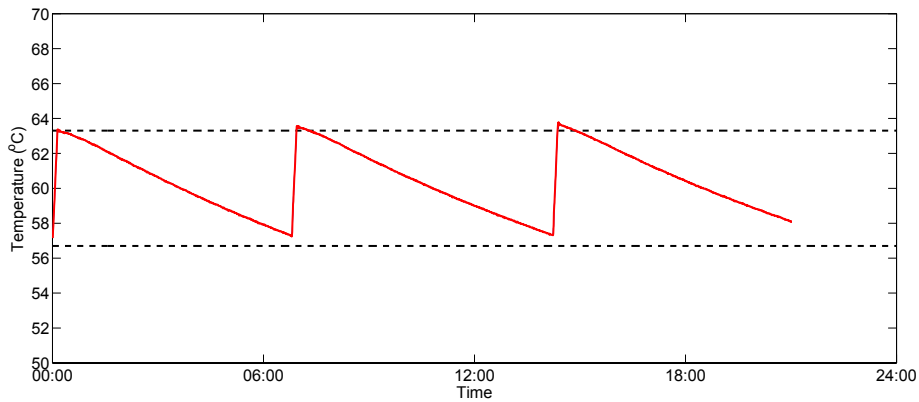
(c) Temperature trajectory of the physical WH.

Figure 6.6. Experiment results of coordinating a population of WHs for renewable integration.

Considering the second case, the solar power deviation and WHs power deviation from baseline are depicted in Figure 6.7 (a). Again, we see the WHs power deviation from baseline tracks the solar generation deviation very well. Additionally, Figure 6.7 (b) shows that the temperatures of simulated WHs are well regulated within the user-specified bounds. However, the WH temperatures get close to the lower temperature bound beginning in the later afternoon, because the WHs are being forced to constantly

6.10

consume less power than the baseline. Their total daily energy consumption is 1743 kWh, which is 92 kWh less than the baseline case of 1835 kWh. Moreover, the average switching time of all WHs during the 24-hour period is 9.5280, which is 3.106 more than the baseline case of 6.1840. Hence, tracking an energy neutral dispatch signal is helpful for reducing the additional ON-OFF cycling times. Nevertheless, for both cases, the WHs exhibit excellent power tracking performance, and the thermal comfort is strictly respected because the temperatures of WHs are well regulated within the user-specified temperature band. Considering the water usage in the model and examining its impacts will be an interesting area of future research.
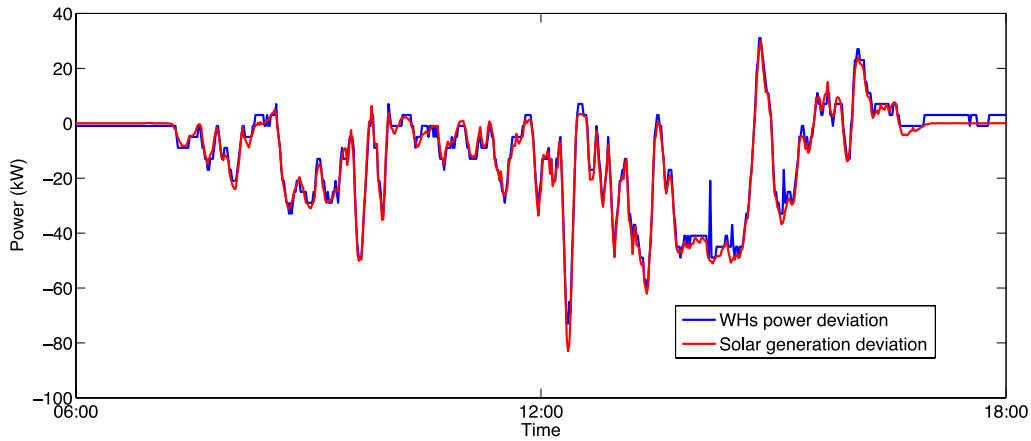

(a) Power-tracking performance
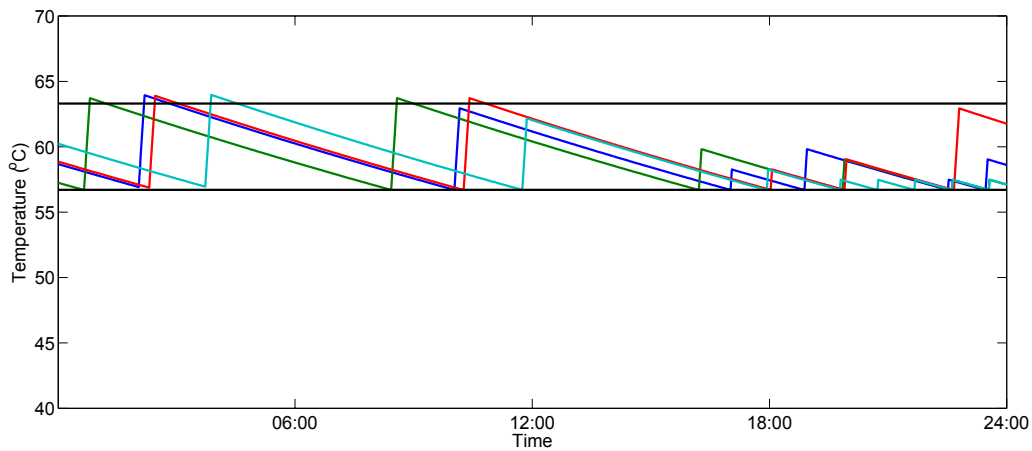
(b) Temperature trajectories of a few simulated WHs.

Figure 6.7. Experiment results of coordinating a population of WHs for renewable integration.

7.0 Conclusions and Future Work

Variable generation sources such as wind and solar power are very uncertain and intermittent. The deep penetration of variable generations into the grid makes power balancing very challenging. Currently, handling the uncertainty of renewables is achieved mostly through supply-side generation reserves. However, this method is very inefficient, from both economic and environmental points of view. A paradigm shift is anticipated to occur in the future, and we must enable demand-side resources to compensate for uncertain generation on the supply side.

In this report, we studied coordination and control of flexible loads for renewable integration. We first presented the motivation and background, and conducted a comprehensive literature review of building-to-grid integration and the studied flexible loads. We then compiled a catalog of flexible building loads, and discussed their characteristics. We next collected solar generation data from a solar PV panel on the PNNL campus and analyzed the data. We found the solar output had a high degree of uncertainty, and the uncertainty occurred at almost all detectable time scales. We also studied the solar generation data from other sources, and verified the results of our previous data analysis. We proposed two transactive coordination strategies to manage flexible building loads for renewable integration. We proved the theories that support the two transactive coordination strategies and discussed the associated tradeoffs between algorithm convergence rate and required amount of information exchange. In our study, we focused on three types of flexible building loads—air handling unit, rooftop unit, and a population of WHs. For each type of load, we presented the system description, model identification, controller design, test bed setup, and demonstration results. We showed that coordination and control of flexible loads had great potential to integrate variable generation sources. The flexible loads could successfully track a power dispatch signal from the coordinator, while having little impact on the quality of service to the end-users.

In our current study, we conducted demonstrations mostly through simulation, and HIL-based studies because of test bed availability. In the future, we would like to implement the demonstrated control strategies on physical test beds, and examine their performance and impacts on the quality of service to the end-users. We are also interested in conducting a live demonstration of the coordination of a collection of flexible loads to absorb the real-time power fluctuations of renewable generation.

8.0   References

ASHRAE. "The ASHRAE Handbook HVAC Systems and Equipment (SI Edition)." 2008.

Bacher, Peder, Henrik Madsen, and Henrik Nielsen. "Online short-term solar power forecasting." *Solar Energy* 83, no. 10 (2009): 1772-1783.

Barooah, Prabir, Ana Buic, and Sean Meyn. "Spectral decomposition of demand-side flexibility for reliable ancillary services in a smart grid." *Hawaii International Conference on System Sciences.* 2015. 2700-2709.

Bashash, Saeid, and Hosam K Fathy. "Modeling and control of aggregate air conditioning loads for robust renewable power management." *IEEE Transactions on Control Systems Technology* 21, no. 4 (2013): 1318-1327.

Boyd, Stephen, and Lieven Vandenberghe. *Convex optimization.* Cambridge university press, 2004.

Brzozowski, Carol. "The "Smart" Water Grid." *Water Efficiency*, 2011.

Building Energy Data Book. http://buildingsdatabook.eren.doe.gov/default.aspx .

California Energy Commission's Public Interest Energy Research Program . "Lighting the Way to Demand Response Lighting the Way to Demand Response." 2011. https://www.esource.com/system/files/files/2011-04/CEC-TB-47_Lighting_web_7.pdf.

Callaway, Duncan S. "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy." *Energy Conversion and Management* 50, no. 5 (2009): 1389-1400.

Chen, Changsong, Shanxu Duan, Tao Cai, and Bangyin Liu. "Online 24-h solar power forecasting based on weather type classification using artificial neural network." *Solar Energy* 85, no. 11 (2011): 2856-2870.

Chen, Lijun, Na Li, Steven H Low, and John C Doyle. "Two market models for demand response in power networks." *IEEE SmartGridComm.* 2010. 397-402.

Chen, Shiyao, Yuting Ji, and Lang Tong. "Deadline scheduling for large scale charging of electric vehicles with renewable energy." *IEEE Sensor Array and Multichannel Signal Processing Workshop.* 2012. 13-16.

Dassault Systemes. 2016. http://www.3ds.com/products-services/catia/products/dymola/.

Department of Energy. *Agricultural demand response program in california helps farmers reduce peak electricity usage, operate more efficiently year-round.* http://energy.gov/sites/prod/files/Case%20Study%20-%20M2M%20-%20Agricultural%20Demand%20Response%20Program%20Helps%20CA%20Farmers%20-%20January%202012.pdf (accessed June 2016).

—. "EnergyPlus engineering reference." 2010. https://energyplus.net/sites/default/files/pdfs_v8.3.0/EngineeringReference.pdf.

DOE Global Energy Storage Database. http://www.energystorageexchange.org/projects (accessed June 2016).

Energy Information Administration. "Commercial Buildings Energy Consumption Survey 2012." July 2016. https://www.eia.gov/consumption/commercial/.

Ghatikar, Girish, and at al. "Demand response opportunities and enabling technologies for data centers: Findings from field studies." Lawrence Berkeley National Laboratory, 2012.

Hao, He, and Wei Chen. "Characterizing flexibility of an aggregation of deferrable loads." *IEEE Conference on Decision and Control.* 2014. 4059-4064.

Hao, He, Borhan M Sanandaji, Kameshwar Poolla, and Tyrone L Vincent. "A generalized battery model of a collection of thermostatically controlled loads for providing ancillary service." *Annual Allerton Conference on Communication, Control, and Computing.* 2013. 551-558.

Hao, He, Borhan M Sanandaji, Kameshwar Poolla, and Tyrone L Vincent. "Aggregate flexibility of thermostatically controlled loads." *IEEE Transactions on Power Systems* 30, no. 1 (2015): 189-198.

Hao, He, Jianming Lian, Karanjit Kalsi, and Jakob Stoustrup. "Distributed flexibility characterization and resource allocation for multi-zone commercial buildings in the smart grid." *IEEE Conference on Decision and Control.* 2015. 3161-3168.

Hao, He, Yashen Lin, Anupama Kowli, Prabir Barooah, and Sean Meyn. "Ancillary service to the grid through control of fans in commercial building HVAC systems." *IEEE Transactions on Smart Grid* 5, no. 4 (2014): 2066-2074.

Helman, Udi. "Resource and transmission planning to achieve a 33% RPS in California–ISO modeling tools and planning framework." *FERC Technical Conference on Planning Models and Software.* 2010.

Hirsch, A., J. Clark, M. Deru, K. Trenbath, I. Doebber, and D. Studer. "Pilot Testing of Commercial Refrigeration-Based Demand." National Renewable Energy Laboratory, 2015.

Ihara, Satoru, and Fred Schweppe. "Physically based modeling of cold load pickup." *IEEE Transactions on Power Apparatus and Systems* 9, no. PAS-100 (1981): 4142-4150.

International Energy Agency. "2014 Key World Energy Statistics." International Energy Agency, 2014.

Kiliccote, Sila, Mary Ann Piette, Johanna Mathieu, and Kristen Parrish. "Findings from seven years of field performance data for automated demand response in commercial buildings." Lawrence Berkeley National Laboratory, 2010.

Kunczynski, Yan, and Ben Burger. "Demand Response Desalination." 2014. http://sisyan.com/pdf/WHITE%20PAPER.pdf.

Kundu, Soumya, Nikolai Sinitsyn, Scott Backhaus, and Ian Hiskens. "Modeling and control of thermostatically controlled loads." *the 17-th Power Systems Computation Conference.* 2011.

Li, Na, Lijun Chen, and Steven H Low. "Optimal demand response based on utility maximization in power networks." *IEEE Power and Energy Society General Meeting.* 2011. 1-8.

Li, Sen, Wei Zhang, Jianming Lian, and Karanjit Kalsi. "Market-Based Coordination of Thermostatically Controlled Loads Part I: A Mechanism Design Formulation." *IEEE Transactions on Power Systems* 32, no. 2 (2016): 1170-1178.

Lin, Yashen and Barooah, Prabir and Mathieu, Johanna L. "Ancillary services to the grid from commercial buildings through demand scheduling and control." *American Control Conference.* 2015. 3007-3012.

Lin, Yashen, Prabir Barooah, and Sean P Meyn. "Low-frequency power-grid ancillary services from commercial building HVAC systems." *IEEE SmartGridComm.* 2013. 169-174.

Lin, Yashen, Prabir Barooah, Sean Meyn, and Timothy Middelkoop. "Experimental Evaluation of Frequency Regulation From Commercial Building HVAC Systems." *IEEE Transactions on Smart Grid* 6, no. 2 (2015): 776-783.

Lorenz, Elke, Johannes Hurka, Detlev Heinemann, and Hans Georg Beyer. "Irradiance forecasting for the power prediction of grid-connected photovoltaic systems." *IEEE Journal of selected topics in applied earth observations and remote sensing* 2, no. 1 (2009): 2-10.

Lu, Ning, and David P Chassin. "A state-queueing model of thermostatically controlled appliances." *IEEE Transactions on Power Systems* 19, no. 3 (2004): 1666-1673.

Maasoumy, Mehdi, Jorge Ortiz, David Culler, and Alberto Sangiovanni-Vincentelli. "Flexibility of commercial building hvac fan as ancillary service for smart grid." *arXiv.com.* 2013. arXiv preprint arXiv:1311.6094.

Mai, Weijie, and CY Chung. "Economic MPC of aggregating commercial buildings for providing flexible power reserve." *IEEE Transactions on Power Systems* 30, no. 5 (2015): 2685-2694.

Makarov, Yuri V, Clyde Loutan, Jian Ma, and Phillip De Mello. "Operational impacts of wind generation on California power systems." *IEEE Transactions on Power Systems* 24, no. 2 (2009): 1039-1050.

Malhame, Roland, and Chee-Yee Chong. "Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system." *IEEE Transactions on Automatic Control* 30, no. 9 (1985): 854-860.

Mathieu, Johanna L, Maryam Kamgarpour, John Lygeros, and Duncan S Callaway. "Energy arbitrage with thermostatically controlled loads." *European Control Conference.* 2013. 2519-2526.

Mathieu, Johanna L, Stephan Koch, and Duncan S Callaway. "State estimation and control of electric loads to manage real-time energy imbalance." *IEEE Transactions on Power Systems* 28, no. 1 (2013): 430-440.

Meyn, Sean P, Prabir Barooah, Ana Busic, Yue Chen, and Jordan Ehren. "Ancillary service to the grid using intelligent deferrable loads." *IEEE Transactions on Automatic Control* 60, no. 11 (2015): 2847-2862.

Meyn, Sean, Prabir Barooah, Ana Busic, and Jordan Ehren. "Ancillary service to the grid from deferrable loads: the case for intelligent pool pumps in Florida." *IEEE Conference on Decision and Control.* 2013. 6946-6953.

Modelica Association. *Modelica-A Unified Object-Oriented Language for Physical Systems Modeling-Language Specification Version 3.3 Revision 1, 2014.* 2014. http://www.modelica.org.

Motegi, Naoya, Mary Ann Piette, David S Watson, Sila Kiliccote, and Peng Xu. "Introduction to commercial building control strategies and techniques for demand response." Lawrence Berkeley National Laboratory, 2007.

NASA Goddard Institute for Space Studies. Goddard Institute for Space Studies (accessed June 2016).

Nayyar, Ashutosh, Josh Taylor, Anand Subramanian, Kameshwar Poolla, and Pravin Varaiya. "Aggregate Flexibility of a Collection of Loads." *IEEE Conference on Decision and Control.* 2013. 5600-5607.

Newquay Weather Station. "Solar PV generation." http://www.newquayweather.com/wxsolarpv2015.php (accessed June 2016).

Nutaro, James, David Fugate, Teja Kuruganti, Jibonananda Sanyal, and Michael Starke. "Cost-effective retrofit technology for reducing peak power demand in small and medium commercial buildings." *Science and Technology for the Built Environment* 21, no. 6 (2015): 761-772.

Otter, Martin, Martin Malmheden, Hilding Elmqvist, Sven Erik Mattsson, and Charlotta Johnsson. "A New Formalism for Modeling of Reactive and Hybrid Systems." *the 7th Modelica Conference.* 2009. 364-377.

Pacific Northwest National Laboratory. *VOLTTRON™.* July 2016. http://transactionalnetwork.pnnl.gov/volttron.stm.

Papavasiliou, Anthony, and Shmuel S Oren. "Supplying renewable energy to deferrable loads: Algorithms and economic analysis." *IEEE PES General Meeting.* 2010. 1-8.

Smith, J. Charles, Michael R Milligan, DeMeo A Edgar, and Brian Parsons. "Utility Wind Integration and Operating Impact State of the Art." *IEEE Transactions on Power Systems* 22, no. 3 (August 2007): 900 - 908.

Stocker, Thomas F., Dahe Qin, Gian-Kasper Plattner, and et al. "Technical Summary." *Climate Change 2013: The Physical Science Basis*, 2014: 33-115.

Subramanian, Anand, Manuel J Garcia, Duncan S Callaway, Kameshwar Poolla, and Pravin Varaiya. "Real-time scheduling of distributed resources." *IEEE Transactions on Smart Grid* 4, no. 4 (2013): 2122-2130.

Todd, D., M. Caufield, B. Helms, M. Starke, B. Kirby, and J. Kueck. "Providing Reliability Services through Demand Response: A Preliminary Evaluation of the Demand Response Capabilities of Alcoa Inc." 2009.

Vrettos, Evangelos, Emre C Kara, Jason MacDonald, Goran Andersson, and Duncan S Callaway. "Experimental Demonstration of Frequency Regulation by Commercial Buildings-Part I: Modeling and Hierarchical Control Design." *arXiv preprint arXiv:1605.05835*, 2016.

Vrettos, Evangelos, Emre C Kara, Jason MacDonald, Goran Andersson, and Duncan S Callaway. "Experimental Demonstration of Frequency Regulation by Commercial Buildings-Part II: Results and Performance Evaluation." *arXiv preprint arXiv:1605.05558*, 2016.

Wang, Weimin, Srinivas Katipamula, Hung Ngo, Ron Underhill, Danny Taasevigen, and Robert Lutes. "Advanced Rooftop Control (ARC) Retrofit: Field-Test Results." Pacific Northwest National Laboratory, Richland, 2013.

Weightman, David, and Jenny Field. "Office Plug Loads: Energy Use and Savings Opportunities." 2012. https://www.esource.com/system/files/files/2012-01/ESource-CEC-WC-1-12-PlugLoads_0.pdf.

Wetter, Michael, Wangda Zuo, Thierry Nouidui, and Xiufeng Pang. "Modelica Buildings library." *Journal of Building Perfromance Simulation* 7, no. 4 (2015): 253-270.

Wierman, Adam, Zhenhua Liu, Iris Liu, and Hamed Mohsenian-Rad. "Opportunities and challenges for data center demand response." *International Green Computing Conference*. 2014. 1-10.

Wikipedia. *Ziegler–Nichols method.* July 2016. https://en.wikipedia.org/wiki/Ziegler–Nichols_method.

Yin, Rongxin, Peng Xu, Mary Ann Piette, and Sila Kiliccote. "Study on Auto-DR and pre-cooling of commercial buildings with thermal mass in California." *Energy and Buildings* 42, no. 7 (2010): 967-975.

Zhang, Wei, Jianming Lian, Chin-Yao Chang, and Karanjit Kalsi. "Aggregated modeling and control of air conditioning loads for demand response." *IEEE Transactions on Power Systems* 28, no. 4 (2013): 4655-4664.

8.6

Appendix A

Distribution

  # Name
      Organization
      Address
      City, State and ZIP Code

  # Organization
      Address
      City, State and ZIP Code
          Name
          Name
          Name
          Name
          Name (#)

  # Name
      Organization
      Address
      City, State and ZIP Code

No. of Copies		**No. of Copies**	
#	**Foreign Distribution**	**#**	**Local Distribution**
#	Name		Pacific Northwest National Laboratory
	Organization	Name	Mailstop
	Address	Name	Mailstop
	Address line 2	Name	Mailstop
	COUNTRY	Name	Mailstop
		Name	(PDF)