Agent-Based System for Transactive Control of Smart Residential Neighborhoods

M. Starke, J. Munk, H. Zandi, T. Kuruganti, H. Buckberry, J. Hall
Oak Ridge National Laboratory
Oak Ridge, TN 37831
{starkemr, munkjd, zandih, kurugantipv, buckberryhl, halljr}@ornl.gov

J. Leverette,
Southern Company
Birmingham, AL, USA
jalevere@southernco.com

***Abstract*— Buildings consume 74% of the total electricity produced in the United States. A significant portion of the building electric load includes heating ventilation and air-conditioning (HVAC) systems and water heating (WH) systems. Enabling flexibility in the operations can improve overall electric grid efficiency. This paper describes a multi-agent system for supporting integration, learning, optimization, and control of HVAC and WH in supporting a future smart grid. The architecture supports a transactive-based negotiation strategy between homeowners and a microgrid controller to adjust consumption behavior and reduce electricity costs. The framework is deployed in a neighborhood and preliminary testing is underway. The agent architecture design is discussed along with the preliminary optimization results from the demonstration site.**

***Index Terms*-- IoT, agents, transactive, demand management, grid-interactive**

I. Introduction

In recent years, electric distribution system operations have been advancing rapidly, particularly with the growing intelligence at the residential level. The Internet of Things (IoT) has brought opportunities to modify setpoints on thermostats and turn on/off devices manually or through a schedule. While load shedding control options from a centralized system have been researched significantly in the past [1]-[5] opportunities to actively control load through a negotiation process has not been researched as much [6]. For example, in [7] a multi-agent energy management solution is proposed that uses central coordination and building management to achieve user objectives, renewable energy source forecasting, and battery bank management. In [8], a fuzzy-based multi-agent centralized energy management system is proposed where each agent receives the measured quantities of the microgrid as input signals. A particle swarm framework is used in [9] as the basis for a facilitator agent which enables neighborhood optimization while building level agents to manage building level optimization. The pilots [10] and [11] use bids, from a satisfaction-based bid curve, from the end-users in the process of setting real-time prices.

This paper focuses on a research and demonstration project with an intelligent negotiation strategy between residential homeowners and a microgrid controller. This strategy is supported by a multi-agent system infrastructure that learns, optimizes, and controls the residential loads. The objective is to develop a robust formulation for transactive control and management of residential loads to reduce energy cost for both homeowners and the utility and maximize the use of microgrid resources. The belief is that by providing a cost incentive to the homeowner to adjust consumption, the overall load profile can be modulated to support utility-scale needs such as peak reduction and renewable energy penetration. Historically, the entirety of the decision-making process has been performed by the utility or microgrid. For example, real-time prices are set by the utility without input from the end-user [12] and [13]. This project is investigating a compromise on decision and control over loads by demonstrating a negotiation-like process between the homes and microgrid or utility. The concept and initial optimization results are presented.

II. Background

A community composed of sixty-two detached, single-family homes and supported by a microgrid has been constructed. An aerial view of the community is shown in Fig. 1. Each home has connected equipment and devices that provide advanced control features and data. For negotiation purposes, the heating, ventilation, and air conditioning systems (HVAC) and water heaters are the primary loads that are engaged for providing flexibility and control. Each of the homeowners has entered into a two-year research agreement with the utility company to allow data collection as well as dispatch control of their HVAC and water heater. The research project included an eighteen-month development phase as the homes were designed and constructed. The research analysis phase will extend another six months to a year after the data collection phase.

The microgrid was constructed to support the community through improved resilience. When a distribution system anomaly is detected, such as an outage, the microgrid islands the neighborhood and continually provides power to the community. The microgrid assets include large-scale solar, energy storage, and a natural gas generator. A microgrid controller has been programmed to optimally determine the generation resource dispatch as well as the potential to incentivize homeowners to adjust electric power consumption. The agent architecture and negotiation process are described in the next section.


Figure 1. Aerial view of Reynold's Landing at Ross Bridge, Smart NeighborhoodTM. Image: Southern Co.

III. AGENT FRAMEWORK AND STRATEGY

The approach for the negotiation process for this project has been developed with the concept of price signal and load curve exchanges. The negotiation process is shown in Fig. 2. The microgrid controller is responsible for issuing a 24-hour projected price signal on a per phase basis (AN/BN/CN) that is based on optimal projections of local generation costs versus achievable load reductions through load adjustment. This information is passed down to all homes for decision-making with a specific iteration number to link the home response back to the price signal. At the home-level, optimization, models, and learning heuristics provide three 24-hour load forecasts for the (1) HVAC, (2) water heater, and all other uncontrolled loads. The aggregator sums the responses and provides an estimated total load by phase to the microgrid controller. Based on the response, the microgrid controller has the option to modify the price signal or accept the current load projection. This process of sending and receiving price signals and load signals is repeated until the microgrid controller has found an acceptable point of convergence.

A multi-agent system (MAS) architecture was developed to support the required functionality and is shown in Fig. 3 . The utilization of the MAS approach leads to more resilient architecture as the agents are able to self-govern and realize a greater objective. The MAS architecture also supports modularity and can be expanded by simply creating more agents.

The presented architecture utilizes a set of three different systems that are constantly interacting to support the greater purpose of the negotiation process: the aggregator, the single home, and learning. Each of these is represented by a virtual machine actively running within the cloud and used to perform unique functions. Application programming interfaces (APIs), ZMQ message bus and Microsoft Azure Service Bus are the key means for the agents to intercommunicate data throughout the system. The agents were coded in Python.

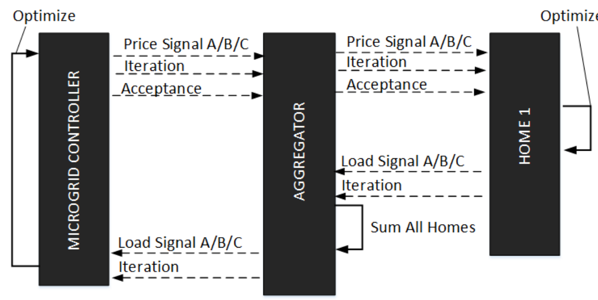

Figure 2. Exchange interaction between Microgrid controller, Aggregator, and home to support negotiation framework.

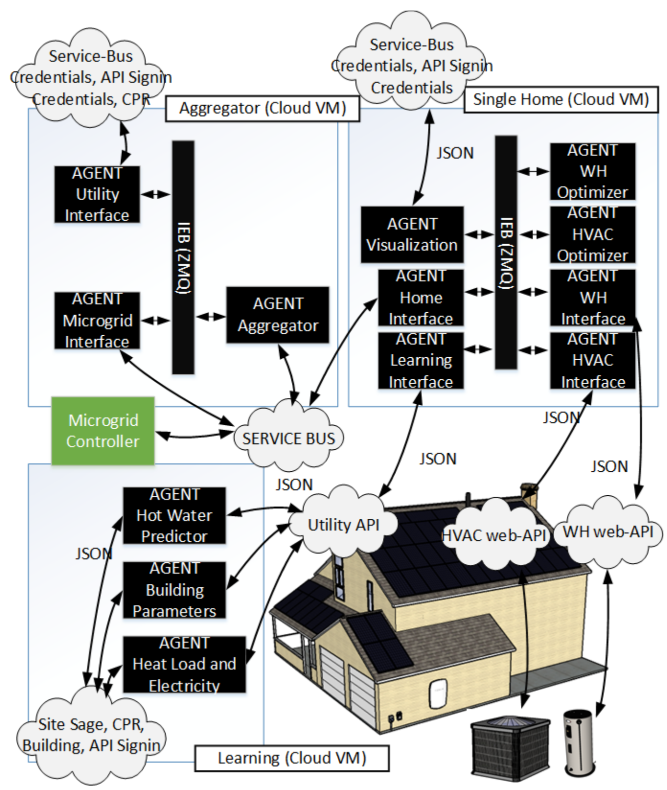

Figure 3. Depiction of agent-based architecture.

The primary purpose of the aggregator system is to distribute and collect information from each of the single home systems within the cloud. The information that needs to be distributed includes credentials and weather forecast from the utility and price signals, acceptance, and iteration from the microgrid controller. The agents within the aggregator system include the utility interface, microgrid interface, and aggregator. The purpose of these agents is described in more detail in TABLE I.

TABLE I. AGGREGATOR INSTANCE AGENTS

Agents	***Purpose***
Utility Interface	Pulls data from Southern Company API which includes weather and Service Bus credentials
Microgrid Interface	Communicates to microgrid controller predicted load consumption, iteration number, and receives price signal, iteration number, and acceptance utilizing Service Bus
Aggregator	Sets up communications to each Single Home VM and issues weather data, price signal, and iteration while receiving predicted load and iteration.

The learning system has the primary purpose of providing predictive elements to the single home system. The learning system extracts recorded data from a historian-based API and attempts to develop and predict several key parameters that are needed for the HVAC and water heater optimization including hot water usage, internal heat load and uncontrolled electrical load, and building model parameters. These agents within the learning system are the hot water predictor, internal heat load and electricity consumption, and building model parameters. The purpose of these agents is presented in TABLE II.

TABLE II. LEARNING INSTANCE AGENTS

Agents	***Purpose***
Hot Water Predictor	Produces a forecast estimation of the hot water usage in time based on flow meter data.
Internal Heat Load and Electricity Consumption	Produces a forecast estimation of heat load within the home and other load electricity consumption based on metered circuit level data.
Building Model Parameters	Estimates the building parameters for the MPC formulation based on historical weather data and HVAC, WH and other metered circuit level data.

The primary system utilized to support decision making is the single home system. The single home system is deployed on identical virtual machines but with unique credentials for each home within the neighborhood. This is necessary as each homeowner has a separate set of log-in credentials for the devices, user inputs for bounding the control, building design, and usage patterns in terms of behavior. These parameters are needed to both communicate to the devices through API calls and to optimize the HVAC and water heater utilization.

The optimization is based on a model predictive control (MPC) approach where a building model and HVAC performance data is used to optimize the HVAC while a water heater model is used to optimize the water heater. The models utilized for the building and water heater are presented in Fig. 4 and Fig. 5. Both models have been used in previous simulation environments [10] and are used here in an MPC formulation. The models are based on electrical equivalencies and have been proven to be able to be learned and tuned [12]. The building parameters, in Fig. 4, that are estimated based on the learning approach are Rroof, Cattic, Rattic, Rw, Cw, Cin, Rim, Cim which represent the thermal resistance and capacitance of the roof, attic, walls, indoor air, and the objects (mass) within the building.

The water heater is modeled in similar fashion as shown in Fig. 5. Since the same water heater equipment is utilized throughout the neighborhood, the thermal resistance and capacitance of the water heater were not expected to change and were hard-coded into the optimization without the need of adjustments from a learning algorithm. However, the hot water draw (Qh) was considered a key element providing an accurate MPC formulation. Qw represents the cold, makeup-water supplied to the water heater.

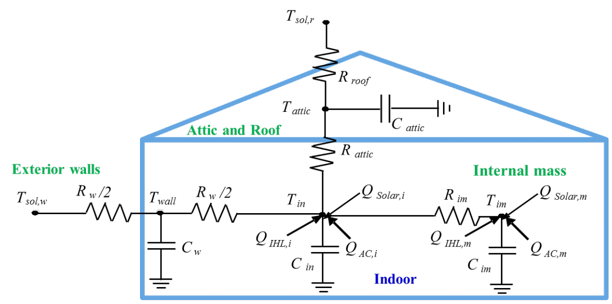

Figure 4. Electrical model representation of residential building used in HVAC optimization.

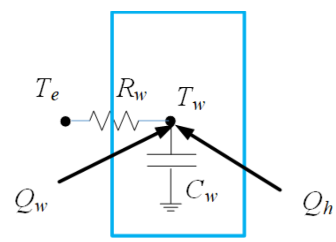

Figure 5. Water heater model used in water heater optimization.

The optimization utilized in the single home system is multi-objective with two primary purposes in mind. The first is to maintain comfort by operating within the homeowner scheduled limits. The comfort is based on temperature scheduled by homeowners on the device. The second is economic and is driven by the price signal issued by the microgrid controller.

The optimization was developed with a linear programming formulation and solved using an open-source solver – computational infrastructure for operations research (COIN|OR) [15] and software interface (PulP) [16]. The linear programming formulation guarantees an optimum solution and the reduced order models solve very quickly on the platform (often within 10 seconds). Comfort was weighted significantly higher than the cost drivers.

The agents within the single home system are the home interface, HVAC interface, Water heater interface, HVAC optimizer, water heater optimizer and learning agent interface. The purpose of each agent is presented in TABLE III.

IV. Implementation

The MAS architecture has been deployed on a system of Microsoft Azure Virtual Machines hosting Ubuntu version 16.04 and VOLTTRON platform. These are hosted in the cloud and managed by the utility with separate collection and automation systems to save each home's HVAC and water optimization results and provide a detailed view into the system behavior. Example results are presented and discussed.

TABLE III. SINGLE HOME INSTANCE AGENTS

Agents	***Purpose***
Home Interface	Data pass-through and collector of optimization and electrical consumption projections for Aggregator agent through Service Bus
HVAC Interface	Gets, Pushes, and Translates HVAC vendor API JSON calls to support decision making and optimization. Heat/Cool Setpoints are primary control means.
Water Heater Interface	Gets, Pushes, and Translates WH vendor API JSON calls to support decision making and optimization. Heat Setpoint and Water Heater Mode are primary control mechanisms.
HVAC Optimizer	Utilizes building specifications, forecasted weather data, building parameter data, price forecast, and HVAC status data to optimally schedule HVAC and provide expected electrical consumption. Issues setpoint.
Water Heater Optimizer	Utilizes predicted water consumption, price forecast, and Water Heater status data to optimally schedule Water Heater and provide expected electrical consumption. Issues setpoint.
Learning Agent Interface	Extracts data from an API to provide input to the optimization models and total load forecast.

Today, initial testing is underway with a baseline price signal based on Alabama Power Real-Time Time of Use Rate for Summer. This price signal is fundamentally a square wave with peak price at $0.21 for hours between 1pm and 7pm local time (or 18:00 and 24:00 UTC) and $0.075 during non-peak hours. The price signal is sent repeatedly for every iteration (adjusted in time based on the microgrid controller clock) for the initial testing. This occurs on a roughly 10-minute basis as the optimization for each reports back in approximately 3 minutes.

Two sets of water heater optimization and control results are presented for two different homes in the neighborhood shown in Fig. 6 and Fig. 7. On the top graphs, the actual temperature measured versus forecasted temperature are shown while in the bottom graphs, the projected electrical consumption versus actual measured is presented. For these homes, the water heater is a hybrid hot water heater with both heat pump and electric operational modes. The key for optimization is to limit electric mode utilization while avoiding high price periods.

In this stage of the work, a generic hot water usage pattern has been implemented for all homes. Since hot water usage is the primary driver for water heater energy use and tank temperature changes, the current model is not able to accurately predict future use or temperatures. This will be addressed in the future by using machine learning to create unique forecasts for hot water usage for each house. Regardless of the poorly matched hot water use forecast, the optimization raised the hot water temperature in both water heaters to ride through both the critical peak price and forecasted hot water draw period. While the actual measured temperature did not identically match projected temperature the optimization and control did successfully preheat the water heater and ride through the critical price period without any utilization.

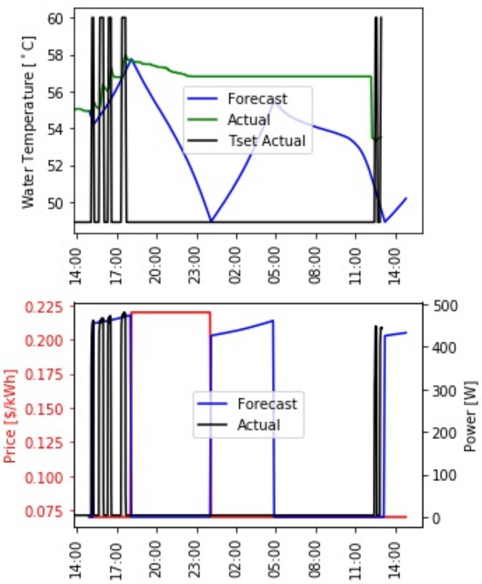

Figure 6. Example comparison of forecasted baseline operation and actual operation for water heater temperature (top) and power consumption (bottom) versus price signal (House A) plotted based at coordinated universal time (UTC)

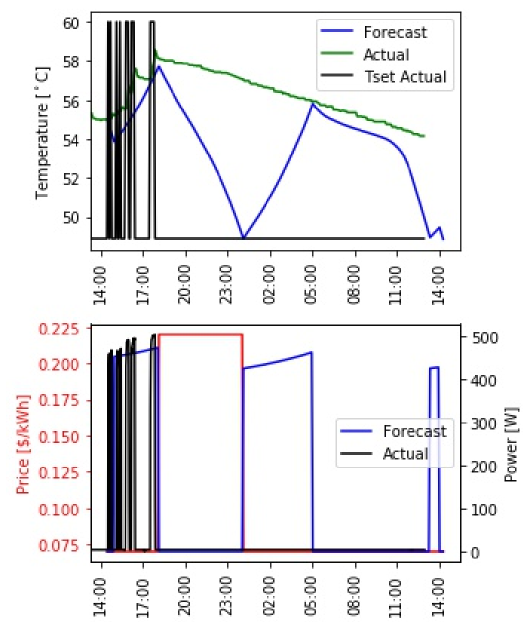

Figure 7. Example comparison of forecasted baseline operation and actual operation for water heater temperature (top) and power consumption (bottom) versus price signal (House B) plotted with UTC.

A single home optimization and control results are presented in Figure 8. In the top graphs, the actual temperature measured versus forecasted temperature are shown while in the bottom graphs, the projected electrical consumption versus actual measured is presented. Similar to the case of the water heater, internal heat load projections have not been tuned to the homes. Additionally, homeowners have the capability to override the temperature at the thermostat. Despite having a programmed temperature range down to 19C in their thermostat, this homeowner overrode the dispatched temperature preventing the home from cooling below 20C.

This resulted in the air conditioner cycling during the peak period contrary to the forecasted results. Additional tuning of the models and communication with the homeowners on how to properly set their desired temperature limits are expected to improve the accuracy of the optimization. This is currently in deployment and early testing.

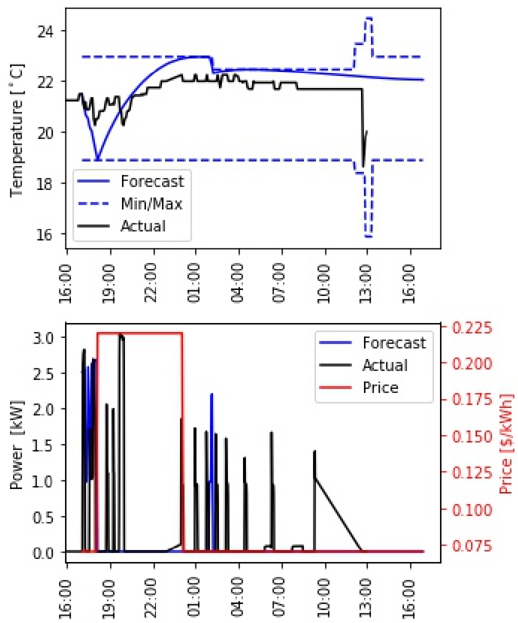

Figure 8. Cumulative cost-optimized power consumption for 10 homes under two different electricity price signals

V. Challenges and Future Work

Some of the challenges encountered during the development and initial deployment (assessment) phases of this project included: interfacing with device APIs where the manufacturers updated the content and/or format of the data, maximum allowable calls to the vendor APIs set to once every 5 minutes, recurring connectivity issues caused by either the homeowners' wi-fi routers or device manufacturer's servers being offline, homeowners electing to manually override their devices rather than changing their programmed setpoints and device point control limitations set by the manufacturer.

Future work includes further assessment of an array of use cases of interest to the utility with respect to home load control or microgrid resource allocation. Further refinements in the optimization algorithms will likely be developed, although much of the changes to the algorithms are expected to occur at the individual house level as a result of the machine learning agent that has been overlaid on the initial agent deployment. Last, homeowner education on input options to devices should reduce override challenges.

V. Conclusion

This paper describes a multi-agent system to support integration and optimization of HVAC and water heaters within a transactive microgrid. The agent system was deployed in a cloud-based system for several weeks and has been controlling residents water heaters and HVAC systems. The system has been collecting performance and optimization results which continue to be evaluated. Data from initial optimization results are shown for a sampling with a real-time price signal. Work continues as evaluation of the ability to load shape and driving factors are assessed.

Acknowledgment

This work was funded by the U.S. Department of Energy, Energy Efficiency and Renewable Energy, Building Technology Office under contract number DE-AC05-00OR22725.

References

[1] D. K. Subramanian, "Optimum Load Shedding Through Programming Techniques," IEEE Trans. Power Appar. Syst., vol. PAS-90, no. 1, pp. 89–95, 1971.

[2] P. Wang and R. Billinton, "Optimum load-shedding technique to reduce the total customer interruption cost in a distribution system," IEEE Proc. - Gener. Transm. Distrib., vol. 147, no. 1, pp. 51–56, 2000.

[3] J. A. Laghari, H. Mokhlis, A. H. A. Bakar, and H. Mohamad, "Application of computational intelligence techniques for load shedding in power systems: A review," Energy Convers. Manag., vol. 75, pp. 130–140, 2013.

[4] M. T. Devine and V. Bertsch, "Examining the benefits of load shedding strategies using a rolling-horizon stochastic mixed complementarity equilibrium model," Eur. J. Oper. Res., vol. 267, no. 2, pp. 643–658, 2018.

[5] U. Rudez and R. Mihalic, "A novel approach to underfrequency load shedding," Electr. Power Syst. Res., vol. 81, no. 2, pp. 636–643, 2011.

[6] S. Chen and C.-C. Liu, "From demand response to transactive energy: state of the art," J. Mod. Power Syst. Clean Energy, vol. 5, no. 1, pp. 10–19, 2017.

[7] A. Anvari-Moghaddam, A. Rahimi-Kian, M. S. Mirian, and J. M. Guerrero, "A multi-agent based energy management solution for integrated buildings and microgrid system," Appl. Energy, vol. 203, pp. 41–56, 2017.

[8] P. Kofinas, A. I. Dounis, and G. A. Vouros, "Fuzzy Q-Learning for multi-agent decentralized energy management in microgrids," Appl. Energy, vol. 219, pp. 53–67, 2018.

[9] M. Hu, J. D. Weir, and T. Wu, "An augmented multi-objective particle swarm optimizer for building cluster operation decisions," Appl. Soft Comput., vol. 25, pp. 347–359, 2014.

[10] D. J. Hammerstrom et al., "Pacific Northwest GridWiseTM Testbed Demonstration Projects; Part I. Olympic Peninsula Project," Pacific Northwest National Lab.(PNNL), Richland, WA (United States), 2008.

[11] S. E. Widergren et al., "AEP Ohio gridSMART demonstration project real-time pricing demonstration analysis," PNNL Rep., vol. 23192, 2014

[12] B. Johnson, M. Starke, O. Abdelaziz, R. Jackson, and L. Tolbert, "A MATLAB Based Occupant Driven Dynamic Model for Predicting Residential Power Demand," in Transmission & Distribution Conference & Exposition (T&D 2014), Chicago, IL, Apr. 14-17, 2014

[13] B.R. Cui, J. Munk, R. Jackson, D. Fugate, M. Starke. 2017. Building thermal model development of typical house in U.S. for virtual storage control of aggregated building loads based on limited available information. The proceeding of ECOS 2017- The 30th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems, July 2 -6, 2017. San Diego, California, USA

[14] P. L. Joskow and C. D. Wolfram, "Dynamic pricing of electricity," Am. Econ. Rev., vol. 102, no. 3, pp. 381–385, 2012.

[15] H. Allcott, "Rethinking real-time electricity pricing," Resour. Energy Econ., vol. 33, no. 4, pp. 820–842, 2011.

[16] COIN-OR available online at https://projects.coin-or.org/Clp, downloaded May 2018.

[17] Pulp available online at: https://www.coin-or.org/PuLP/, downloaded May 2018.