Achieving Controllability of Plug-in Electric Vehicles

Ian Hiskens
Electrical Engineering and Computer Science
University of Michigan, Ann Arbor
hiskens@umich.edu

Duncan Callaway
Energy and Resources Group and Mechanical Engineering
University of California, Berkeley
dcal@berkeley.edu

***Abstract*— The paper presents a conceptual framework for actively involving highly distributed loads in system control actions. The context for load control is established by first reviewing system control objectives, including economic dispatch, automatic generation control and spinning reserve. Also, prior work on load control is reviewed. The load control strategy discussed in the paper builds on the concept of a load aggregator. The aggregator acquires data from plug-in electric vehicle loads in its area, and builds a consolidated model that describes overall load availability. When control actions are required, the aggregator broadcasts a common message to all loads, with the response of individual loads dependent upon their interpretation of that message. The interface between the aggregator and the system controller should have a form that allows load control to be integrated seamlessly into the legacy system. The paper discusses the communications infrastructure required to support such a load control scheme.**

I. Introduction

The purpose of this paper is to explore ways in which plug-in electric vehicles (PEVs) might be controlled for the benefit of power system operating cost, efficiency and reliability. Although PEVs are not yet in mass production (and therefore not yet available for grid services – we will define these system services in Section II), we can look to the history of electricity load control to inform our perspective on PEVs.

In principle, practically any measure that can be taken on the supply side to ensure that electricity generation and load are equal has an equivalent countermeasure on the demand side. However, as we will discuss in this paper, customer expectations for end-use function, as well as limitations in the communications and control infrastructure, influence the degree to which these applications can be served via load control. As a consequence, electricity supply has historically been the key controlled variable for managing power system operation. There are nonetheless several key advantages that may follow from using loads for system services:

1) Although individual loads may become unavailable at any moment, the variability of the total contribution of a very large number of small loads is likely to be less than that of a small number of large generators (for which the failure of one can have substantial impact on the ability to provide the desired service).
2) Loads can often respond to operator requests instantaneously, whereas generators require at least several minutes to make output changes of any significance.
3) Because loads are distributed throughout the grid, they provide the opportunity to devise spatially precise responses to contingencies.
4) In some situations, using loads to provide system services could reduce overall grid emissions (especially if the replaced generation is relatively inefficient).
5) The level of spatial and temporal flexibility that loads could provide to the power system might be used to support the growing penetration of intermittent renewable electricity generators.
6) Loads are already embedded in the power system, and advanced metering infrastructure is bringing communications capability to these loads. Therefore it may soon be the case that all that stands between reliable utilization of loads for system services is the development of the necessary load models and control strategies.

The paper explores conceptual load and PEV control frameworks that are non-disruptive, in the sense that end-use function is not significantly compromised. Section II reviews the current "supply-side" focused grid operating paradigm, while Section III describes some of the existing work on load control. A framework for PEV load control is developed in Section IV, and the related communications requirements are explored further in Section V. Conclusions are presented in Section VI.

II. Conventional power system operation and control

A. Economic dispatch and unit commitment

Electricity demand in a power system varies throughout the day, following patterns that depend on, among other things, regional characteristics, temperature, time of day, day of week, and season of the year. Decisions to change generator output to accommodate variation on hourly time scales are usually made by processes of *unit commitment* (UC) and *economic dispatch* (ED). UC establishes generator operating schedules in advance of the operating time and takes into account generator ramping capabilities and startup and shutdown costs. UC must be solved via a multi-period optimization process such as dynamic programming. Typically UC is carried out one day in advance. ED is the process of choosing generator output levels for available generators to minimize the total cost of meeting demand. ED for committed units can be done within hours or even minutes of the operating time. Consequently both processes require demand forecasts.

This supply-side operating paradigm assumes electricity demand is not controllable or does not change in response

978-1-4244-2601-0/09/$25.00

to changes in supply side operating costs. By and large, this assumption is valid. As a consequence, as electricity demand increases, generation costs can become extremely high (as increasingly inefficient and expensive generation is brought on-line), so much so that supply-side costs exceed the retail price by an order of magnitude or more.

B. Automatic generation control

Although considerable effort goes into predicting electricity demand so that generators can be dispatched as efficiently as possible, it is impossible to predict electricity demand with complete accuracy. Second-to-second and minute-to-minute fluctuations are especially difficult to foresee, and as a result economic dispatch on this time scale can be impractical. Instead, short time-scale variation in electricity demand is usually met by controlling generators that are not necessarily on the economic margin, but are well-suited for the purpose of matching load fluctuations because of their flexibility. A combination of three supply-side mechanisms is typically used to provide this control.

First, an unanticipated change in load or generation is initially compensated for by the addition or extraction of kinetic energy from the rotating inertia of the system's generators; this results in a change in system frequency. Second, many generators are equipped with frequency responsive governors that produce an output change proportional to the frequency deviation (the constant of proportionality is known as a speed-droop characteristic). If system frequency deviates sufficiently far from its setpoint (e.g. 35 mHz or more [14]), droop is activated to prevent further growth of the deviation. This control strategy is inherently decentralized and robust to small disturbances. Furthermore it is initiated almost instantaneously, although a governed generator may require some time to achieve the output level dictated by its droop characteristic. Although feedback control is utilized to achieve the output level dictated by a generator's droop characteristic, there is no feedback on system-level frequency. As a result, generator droop is unable to restore system frequency to its setpoint.

A third mechanism called automatic generation control (AGC) provides feedback control on system-level frequency. (AGC is sometimes referred to as load frequency control or regulation.) AGC decision-making occurs at the level of "balancing authorities" (BAs), which are relatively large regions that might contain hundreds of thousands or millions of customers. When BAs are interconnected, unanticipated changes in load or generation can result in deviations in scheduled interregional tie-line flows as well as frequency deviations. Because both deviations are undesirable, AGC calculations are usually based on a weighted sum of system frequency and unscheduled power flows. The resulting signal is called area control error (ACE). To minimize ACE, AGC issues raise or lower signals based partly on each generator's ability to provide the desired response in a reasonable amount of time and partly on real-time economic dispatch [7]. These signals are typically pulses of varying length (and proportional to the requested output change) that are conveyed on a dedicated communications infrastructure which also telemeters the state of all generators in the BA. Although the signals may be updated based on system ACE and issued as frequently as once every two seconds, economic dispatch targets will not be updated this frequently due to the required computing time.

Unlike the unit commitment problem, AGC need not be a multi-period optimization process (although a model predictive control approach to AGC has been developed in [19]). Instead, decisions can be made based solely on instantaneous generator availability and frequency deviations. However as we will later discuss, engaging loads in the AGC process may require control strategies that forecast how loads will respond to control signals in the future.

C. Contingency reserves

When a sudden, large loss of power supply occurs on the grid (for example a transmission line or generator trips off-line), a large frequency excursion occurs. That causes frequency-responsive generators (referred to as *spinning reserve*) to automatically begin increasing their output to reduce the supply imbalance. Following such an event, it is common for AGC to be disabled until the system operator is able to restore grid frequency (or ACE) to its setpoint by manually issuing raise-lower signals to reserve capacity via the system telemetry infrastructure. This might take 5 to 10 minutes as spinning reserve generators cannot instantaneously increase their output. In order to have sufficient capacity to quickly accommodate a contingency, spinning reserve generators must be grid connected and operating in a part-loaded state. Part-load operation is usually inefficient, so spinning reserve increases operating cost and emissions.

As with AGC, spinning reserve generation need not be dispatched via a multi-period optimization process. Although these generators may be limited in how long they are capable of providing reserve power, the duration of this limitation is typically not binding as system operators can usually bring supplemental reserves online in less than an hour.

III. Prior work on load control

A. Modifying the economic dispatch: load shifting

For several decades, a few utilities have maintained the infrastructure to curtail electricity loads (especially air conditioners and water heaters) to reduce load rather than dispatch additional generation during periods of very high demand. Figure 1 shows a desirable redistibution of load from peak to off-peak. This type of load control is intended to reduce supply side operation costs (by reducing the need to operate high marginal cost peaking generation) and to improve system reliability (by maintaining an acceptable operating reserve). This is one of the few load control applications in use today.

The central challenges associated with controlling electrical loads to contain generation costs are that,

1) The total power and energy available to manage is limited by the obvious need to serve the primary end-use function of the load.
2) There is often a post-control "pickup" in load that results from the continuous operation of previously controlled loads as they recover their desired temperature setpoint. In some circumstances the pickup event can cause total

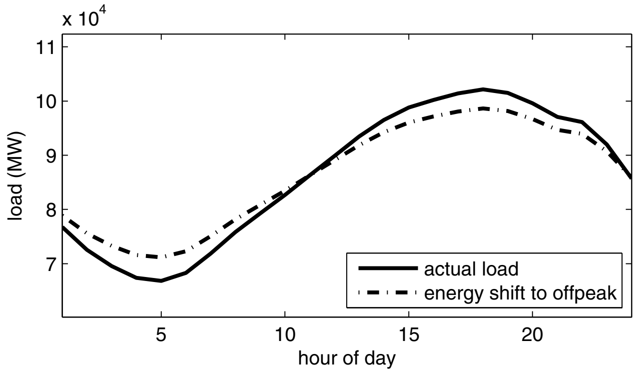

Fig. 1. A hypothetical redistribution of load from peak to off-peak hours. (Original data take from the Midwest Independent System Operator website, www.midwestiso.org.)

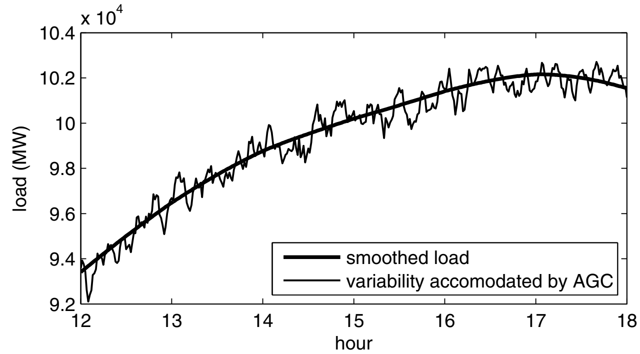

Fig. 2. An example of the rapide time-scale variability that loads or fast-responding generators might in automatic generation control.

load to exceed that which would have occurred in the absence of control.

Some of the most innovative recent advances have dealt with the topic of balancing customer comfort with the need to reduce system load [3], as well as model predictive control approaches to preventing post-peak pickup phenomena [12], [13], [11]. All significant implementations of this type of peak load management curtail loads by engaging relays that interrupt power to the load; these relays are most often activated by a radio signal or by a modulated carrier signal sent directly over the power lines. A few others have explored different control inputs. For example, Navid-Azarbaijani and Banakar [17] have published work on a different control option that adjusts the duty cycle of units via a pulse-width modulated signal. More recently, Burke and Auslander [4] explored adjusting the setpoint of a population of programmable communicating thermostats for peak load shaving.

Today, peak demand management programs that utilize direct load control are disruptive and can have significant impacts on the end-use. In the case of air conditioning, in most regions of the United States load is highest on the warmest days, meaning that air conditioners are curtailed when their services are most in demand. Though the research cited above explores the use of feedback control, actual implementations are open loop and relatively unsophisticated with respect to minimizing impact on the end-use function. This is at least in part due to the historically high cost of reliable sensing equipment; advanced metering infrastructure (AMI) may change this situation.

Plug-in electric vehicles (PEVs) are appealing as controllable loads because they could be curtailed for significant periods of time (several hours) without impact on end-use function. Provided that a vehicle's battery state-of-charge is sufficient by the time it is needed, the vehicle owner has little concern for the details of when and how quickly it is charged. (An exception to this is related to the impact that charging rates may have on battery state of health.) PEVs could be managed not only during peak hours (when their contribution to load may not be very large) but also for night time "valley filling" load control strategies that distribute PEV charging to minimize total energy costs.

B. Managing differences between supply and demand: droop, regulation and energy imbalance

The error between actual load and economically dispatched generation that droop and AGC address is sometimes called energy imbalance. Figure 2 provides an example of what this variability might look like. This error, which is driven by random unpredictable processes, is roughly zero-mean. Therefore electricity loads with some form of capacitance (either thermal in the case of thermostatically controlled loads, or electrical in the case of PEVs) are excellent candidates for imbalance control, especially if the mean error approaches zero over a relatively short averaging time frame.

There are some early stage efforts to manage energy imbalance with frequency-responsive load by providing the equivalent of generator droop [18], [1], [9]. As with generator droop, the approach is completely decentralized, but in the case of loads it may be challenging for system operators to predict the system-level response that will be produced by thousands or millions of unknown devices. There may be more potential in integrating loads into an AGC-type scheme, because it will provide system operator awareness and control of the response. The only effort of which we are aware to provide regulation by some type of centralized control is by one of this paper's authors [5]. However, loads have not been recruited into providing regulation services in any major markets or regions [10]. For several reasons, control will likely need to be distributed or hierarchical. This issue will be discussed in more detail below.

C. System reliability: contingency reserves

Electricity loads are well suited to providing reserves because they can respond very quickly (the ramp rate can be nearly infinite). Indeed, for some time system operators have used non-selective load shedding (i.e. disconnecting entire regions from the grid) as a measure of last resort to avoid system collapse. Selective load shedding (i.e. disconnecting customers or specific customer loads based on pre-arranged agreements), on the other hand, has much more potential from the perspective of customer acceptance. As in the case of using loads to manage energy imbalance, loads with significant capacitance are especially well suited for providing spinning reserve. This is because the time required to restore the system, and relieve loads of their duties, is often on the order of 30 minutes – short enough that the end-use function may not suffer [6]. A number of recent publications and white papers have explored the potential of using responsive loads for spinning reserve [16], [9], [18], demonstration projects are showing promising results[6]. Furthermore, several electricity markets (including ERCOT and PJM in the U.S., and systems in the U.K., Norway and Finland) have instituted programs to use loads as reserves [10].

With advanced metering infrastructure and other emerging grid "cyber-infrastructure" developments, it is becoming increasingly feasible to selectively shed loads based on their individual states. This offers the potential to choose loads whose end-use function will be least affected by the control action. However, as with the services above, there are challenges associated with coordinating thousands or even millions of loads in a way that minimizes end-use impact, or guarantees a certain level of end-use function [2].

IV. PEV LOAD CONTROL

A. Dual objectives of load control

Loads that are controlled to deliver system services must also serve a local control objective (e.g. maintain a thermostat setpoint, deliver a certain PEV state of charge). An effective load control scheme will process these dual objectives. For example, load shifting actions that curtail air conditioners should ideally consider individual building temperatures or occupant comfort levels. This fundamental difference between load control and generator control has several implications, including needs for communications, modeling and control structures that may be challenging to implement. We will discuss these issues in the remainder of this section.

B. General philosophy

The control structure described in Section II plays a fundamental role in the operation of large-scale power systems. That operating paradigm is not likely to change in the foreseeable future, even though some challenges are arising due to the variability of renewable generation. Strategies for incorporating load control into power system operations must therefore ensure consistency with the legacy system. Furthermore, the effectiveness of load control is very much dependent upon participation of large numbers of devices. It is impractical for so many devices to interact directly with high-level power system controls. Load aggregators can serve to satisfy these dual requirements, as shown schematically in Figure 3.

In the remainder of this section, we describe a framework under which load aggregators could engage a very large number of loads within the legacy operating paradigm. We will then use this framework to explore in subsequent sections the requirements for information exchange, load modeling, and controller formulation.

Each load aggregator has jurisdiction over a certain group of loads, and provides an interface between those loads and the higher-level controls. It acquires information from participating loads, describing their ability and willingness to respond

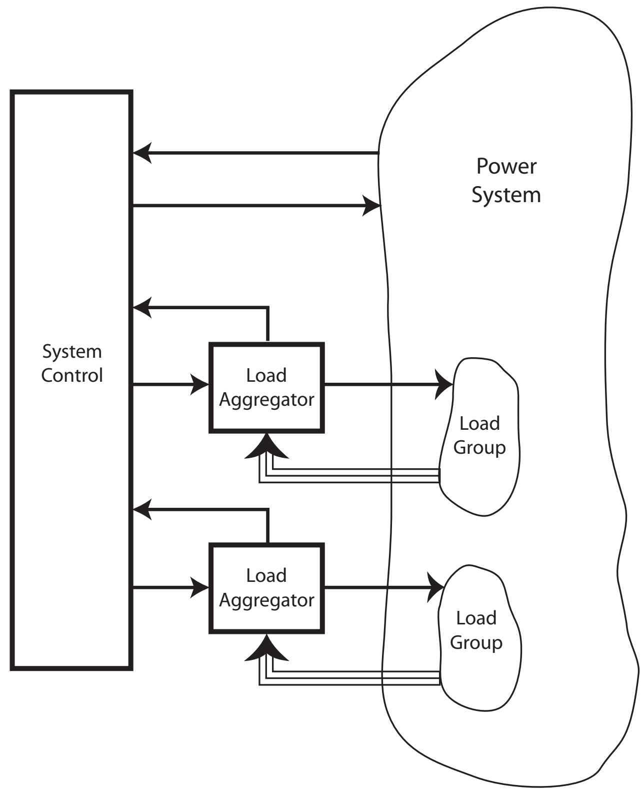

Fig. 3. Schematic representation of a load control strategy.

to control actions. These requirements are discussed further in Section IV-C. The aggregator uses the information provided by individual loads to build a model of the responsiveness of the entire group. The exact form of that model depends on the role that the group may be called upon to perform. For example, in the case of AGC, the model would describe the load increase/decrease achievable in the short term. Model details are discussed further in Section IV-D.

To seamlessly integrate into the existing system, the aggregator should interact with the higher-level controller in a manner consistent with other controllable devices. The two-way information exchange between the controller and aggregator should have the same form as signals from/to generators, for example. Aggregators should be capable of dispatching their loads to respond to the same raise/lower signals that are directed to generators providing automatic generation control. This implies that the aggregator must interpret the control signals received from the higher-level controller, and pass on instructions that are meaningful to the loads.

It cannot be expected that the aggregator will have the capability of tailoring instructions for individual loads. Acquiring information from a large number of loads, for building a consolidated model, is relatively straightforward. It doesn't require a detailed database of individual loads. On the other hand, sending control signals to individual loads would require the aggregator to maintain such a database. Given the mobility of PEVs, the database would need to dynamically update to ensure it accurately reflected the composition of the load group. It is more reasonable to expect the aggregator to broadcast a common signal to all loads in the group, allowing the loads to interpret that signal and respond accordingly.

The exact form of the signal broadcast to loads is an open question. One possibility is to send a value that varies over a predefined range, for example, $+1$ to -1. Here, $+1$ corresponds to switching on all available loads, 0 indicates no change to the pre-existing load level, and -1 provides a directive to switch off all controllable loads. As the load-control signal varied, PEV charging loads would switch on and off as appropriate. It might also be possible to eventually direct PEVs to discharge to the grid in a "vehicle-to-grid" (V2G) mode in response to a similar type of signal [15].

C. Load information requirements

For the load aggregator to build a model of load controllability, it must acquire information describing the availability of load, and its willingness to participate. An illustration may help clarify these concepts. Consider a PEV that is plugged in at 6:00pm, requiring 8 kWh to be fully charged. The owner does not care when charging occurs, providing it is completed by 7:00am the next morning. The maximum charge rate is 2 kW. During the early evening, the charging load is available and willing to be controlled (switched in), though probably would not be called upon. It may be switched on later in the evening, in which case it would be available and willing to be switched off, if necessary. Given that it will take 4 hr to charge, if it has not been switched in as 3:00am approaches (4 hr prior to the deadline), its willingness to participate in control will diminish. Come 3:00am, it will no longer be available for control, as it must switch in to satisfy the charging constraint. The load's availability will always range between ± 2 kW, with the sign dependent upon whether or not it is switched in. Its willingness to participate will be determined by balancing the amount of energy still required and the time remaining.

D. Aggregator load models

Using load availability and willingness measures, the aggregator can build a model that describes load controllability at that point in time. Such a model, however, does not allow the aggregator to best manage temporal constraints. Linking back to the earlier illustration, the aggregator may institute control decisions that leave insufficient controllability as the morning peak approaches. There would seem to be benefits in loads providing the aggregator with information describing their energy requirements and delivery time-frame. The aggregator could then build, and continually update, a load control schedule that maximized controllability over a finite horizon while satisfying energy delivery constraints. The challenge then lies in communicating these constraints to the system operator in a way that leads to control decisions that maximize the use of loads without compromising end use function.

V. Communications requirements

A. Infrastructure

The roll out of advanced metering infrastructure (AMI) offers a communications infrastructure that is suited to communicating load information to aggregators. AMI takes different forms, but typically consists of a home area network that communicates with the electricity meter, a wireless local area network that collects meter information in one "cell relay", and a broadband connection for passing that meter information from the cell relay onto an AMI collection point. For garage-based PEV charging, the PEV would link with the AMI via the home area network. For street-side charging, a link would be required between the supply-point meter and the wireless meter collection network.

It is envisaged that AMI collection points will coincide with distribution substations. Load aggregators may reside at that level, though would more likely draw together load information from a number of AMI collection points. It is important that the total load associated with each aggregator is sufficient for meangingful participation in system control functions. Furthermore, by covering a large number of PEVs, aggregators will be able to achieve relatively smooth variations demand.

As mentioned earlier, the most practical forms of load control tend to utilize return control signals that are broadcast across all loads, rather than targeted to specific installations. Such signals could be delivered via the AMI network, though alternatives include delivery over the internet in conjunction with energy price information. The challenge of the latter option is that the communications network would be owned by another party that is not directly involved in power system operations. This could lead to complications surrounding maintenance, reliability and the ability to issue very high priority signals to the loads under control.

B. Bandwidth requirements

The bandwidth required to support fully functional load control is not significant. Very little information is required to describe each PEV load. However, current AMI wireless LAN communications protocols may significantly limit the rate at which data can be collected. This is because meter read requests are issued at a limited frequency, on the order of once in 5 seconds. If unique requests are required for each unique meter, as is currently the case for collecting AMI electricity consumption data, then it would take one cell relay well over an hour to query 1000 meters (a typical LAN size). The AMI network also needs to collect electricity consumption data for billing, and it typically does this 3 times per day. With these frequency limitations and competing AMI network uses, the ability to collect load state information for control purposes may be challenging. On the other hand, because the update rate is on the order of seconds, if the relay can be configured to collect data from all meters in its LAN once per update, then the speed of information gathering is likely to be more than adequate.

It is envisaged that aggregators will involve relatively simple algorithms, so computational requirements will not be significant. Given that data transfers should be secure, the greatest computational burden may be associated with encryption/decryption algorithms.

C. Data ownership

Data ownership is an issue that needs to be carefully considered in the development of load control schemes. The actual load data will be owned by the distribution company, and will have commercial value. On the other hand, the system control functions will typically be the responsibility of a different organization that has responsibility for operating the overall system. The load aggregator provides an interface between the raw data coming from the loads, and the consolidated model used for system operation and control.

D. Latency

It is well known that time delays within control loops can result in degraded performance and even instability [8]. Time delays in the measurement process cause the controller to operate on old information. A time delay in the actuation process, on the other hand, results in the control action influencing the system later than intended. In both cases, it is very likely that closed-loop performance will be degraded, especially when fast response times are required and/or frequent control updates are issued – as with spinning reserve and automatic generation control.

In the case of load control, both upstream and downstream communications processes are subject to time delays. Upstream data transfer consists of loads communicating with their aggregator, and the aggregators passing consolidated information onto the system controller. As mentioned earlier, AMI is a likely candidate for the first phase of this communications path. The latency associated with AMI technology is not yet well documented, but appears to be on the order of 3 to 8 seconds. This will form the major delay in making load models available to the system controller. The commands generated by the controller could be broadcast over the internet. Latency of internet traffic is uncertain, but is typically quite small.

The load control structure presented in Figure 3 effectively decouples the process of building aggregate load models from the use of those models. As discussed above, most of the communications delay is confined to the model building process. Consequently, the most significant impact of latency occurs through the system controller's use of models that may be up to 10 seconds out of date. This is insignificant under normal load variation conditions. However, if the controller calls for a large load change, for example in response to a need for spinning reserve, the delay in model rebuilding may result in subsequent control actions that are inaccurate and potentially destabilizing. Further work is required to explore these issues.

VI. Summary and Conclusions

This paper has explored the implications of engaging PEVs in system level grid operations such as automatic generation control, spinning reserve and economic dispatch / unit commitment. The central challenge for PEV load control for system services lies in ensuring adequate end-use performance (i.e. state of charge) while also delivering system-level services. Meeting this challenge requires load models, communications and control frameworks that can balance these objectives effectively. This paper's conceptual framework serves as a starting point for understanding these requirements. Understanding communications and load model requirements will be central to moving forward with PEV load control. We have focused on load aggregators as the foundation of the framework, because aggregators have the potential to interact with the system operator in much the same way as generators do. In future work we intend to develop the mathematical structures that might be employed to implement this scheme.

References

[1] Anon. Dynamic demand control of domestic appliances. Technical report, Market Transformation Programme, 2008.

[2] R. Belhomme, R. Cerero Real De Asua, G. Valtorta, A. Paice, F. Bouffard, R. Rooth, and A. Losi. ADDRESS - Active demand for the smart grids of the future. In *Proceedings of the CIRED Seminar 2008: GmartGrids for Distribution*, Paper No 0080, Frankfurt, Germany, June 2008.

[3] K. Bhattacharyya and M.L. Crow. A fuzzy logic based approach to direct load control. *IEEE Transactions on Power Systems*, 11(2):708–714, May 1996.

[4] W. Burke and D. Auslander. Robust control of residential demand response network with low bandwidth input. In *Proceedings of the ASME Dynamic Systems and Control Conference*, Ann Arbor, MI, October 2008.

[5] D.S. Callaway. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management*, 50(9):1389–1400, May 2009.

[6] J.H. Eto, J. Nelson-Hoffman, C. Torres, S. Hirth, B. Yinger, J. Kueck, B. Kirby, C. Bernier, R. Wright, A. Barat, and D.S. Watson. Demand response spinning reserve demonstration. Technical Report LBNL-62761, Lawrence Berkeley National Laboratory, May 2007.

[7] J.D. Glover, M.S. Sarma, and T. Overbye. *Power System Analysis and Design*. CL Engineering, New York, NY, 2007.

[8] G. Goodwin, S. Graebe, and M. Salgado. *Control System Design*. Upper Saddle River, New Jersey: Prentice Hall, 2001.

[9] D.J. Hammerstrom, J. Brous, and T.A. et al Carlon. Pacific Northwest GridWise^{TM} testbed demonstration projects, Part II. GridFriendly^{TM} appliance project. Technical Report PNNL-17079, Pacific Northwest National Laboratory, October 2007.

[10] G. Heffner, C Goldman, B. Kirby, and M. Kintner-Meyer. Loads providing ancillary services: Review of international experience. Technical Report LBNL-62701, Lawrence Berkeley National Laboratory, May 2007.

[11] Y.Y. Hsu and C.C. Su. Dispatch of direct load control using dynamic programming. *IEEE Transactions on Power Systems*, 6(3):1056–1061, Aug 1991.

[12] K.Y. Huang, H.C. Chin, and Y.C. Huang. A model reference adaptive control strategy for interruptible load management. *IEEE Transactions on Power Systems*, 19(1):683–689, Feb. 2004.

[13] K.Y. Huang and Y.C. Huang. Integrating direct load control with interruptible load management to provide instantaneous reserves for ancillary services. *IEEE Transactions on Power Systems*, 19(3):1626–1634, Aug. 2004.

[14] N. Jaleeli, L.S. VanSlyck, D.N. Ewart, L.H. Fink, and A.G. Hoffmann. Understanding automatic generation control. *IEEE Transactions on Power Systems*, 7(3):1106–1122, Aug 1992.

[15] W. Kempton and J. Tomić. Vehicle-to-grid power implementation: From stabilizing the grid to supporting large-scale renewable energy. *Journal of Power Sources*, 144(1):280–294, June 2005.

[16] B. Kirby. Spinning reserve from responsive loads. Technical Report ORNL/TM-2003/19, Oak Ridge National Laboratory, March 2003.

[17] N. Navid-Azarbaijani and M.H. Banakar. Realizing load reduction functions by aperiodic switching of load groups. *IEEE Transactions on Power Systems*, 11(2):721–727, May 1996.

[18] J.A. Short, D.G. Infield, and L.L. Freris. Stabilization of grid frequency through dynamic demand control. *IEEE Transactions on Power Systems*, 22(3):1284–1293, Aug. 2007.

[19] A.N. Venkat, I.A. Hiskens, J.B. Rawlings, and S.J. Wright. Distributed mpc strategies with application to power system automatic generation control. *IEEE Transactions on Control Systems Technology*, 16(6):1192–1206, Nov. 2008.