

SmartLink

A Hierarchical Approach for Connecting Smart Buildings to Smart Grids

Paulo Carreira and Renato Nunes
Dept. of Computer Science and Engineering
Instituto Superior Técnico
Lisboa, Portugal
{paulo.carreira, renato.nunes}@ist.utl.pt

Vasco Amaral
CITI, Dept. of Computer Science
Faculdade de Ciências e Tecnologia, UNL
Caparica, Portugal
Email: vasco.amaral@di.fct.unl.pt,

Abstract—It is widely acknowledged that Automated Demand Response (ADR) is a key factor for the success of the Smart Grid. However, most current ADR implementations are implemented through some form of direct load control, where the utility is handed control to remotely shutdown some agreed electric load at the consumers site. However, being aware of context details local to consumer would lead to actions that potentially achieve greater savings. The role of agents that stand in between the utility and clients known as aggregators have received great attention as they can, in principle, explore this context information. We argue that the current model of DR is too limiting of the consumption reduction potential of ADR and champion a fully decoupled hierarchical ADR model, which is based on a more general definition of aggregator.

In this article we present the concept of an hierarchical ADR model designated as *SmartLink*, where different types of aggregator nodes are connected together forming a tree of decision nodes which are capable of cooperatively create highly adaptable load control strategies to meet a given load reduction targets and taking into account the consumer's local specificities, preferences and utilization constraints.

Keywords—*smart grid; automated demand-response; building automatio; software engineering.*

I. INTRODUCTION

The rigidity of the traditional grid is a major hindrance to the adoption of renewables, to on-site generation as well to the deployment of plug-in electrical vehicles [1]. The concept of a Smart Grid (SG), a computerized power grid that takes advantage of information and communication technology to provide many advanced services has been called to the rescue [2]. However, due to the enormous costs involved in this upgrade, grid operators are looking for ways to leverage the new services provided by the SG to offset the costs. One promising solution seems to lie in a greater participation on the demand side by what is commonly called virtual generation. The fundamental idea is that on critical peak periods, the grid would schedule the consumers to reduce their loads, thus acting as if they had power generating capacity of their own (though some may actually have) and economically compensate them for their participation. The most natural implementation of this concept relies on promoting Demand Response (DR) [3], i.e., in increasing the participation of the consumers by exposing them to dynamic pricing of energy. The premise is that in face

of more precise price signals customers will act with economic rationality [4] to the incentive and engage in (i) performing demand shifting, transferring their electric loads from high price periods to low price off-peak periods or, otherwise (ii) in performing demand reduction by cutting their loads or performing self-generation whenever it makes economic sense to do so. The economic benefits realized from the wide adoption of DR are expected to pay the largest share of the investment on the SG [5].

Despite the fact that consumers do react to price variations [6], the magnitude of the response for both domestic and industrial customers is highly variable. It has been found that the most responsive consumers depend on equipment that assists them in DR [7], [8], [9] and that the need for human intervention is hampering DR [8]. Conceivably, the planning overhead and effort required to ensure business continuity while taking advantage of dynamic pricing is simply infeasible for many customers. Hence, the development of automated DR (ADR) solutions that free the consumer from the concerns of performing demand shifting and demand reduction by controlling electric loads automatically, and in price-responsive fashion, is expected to be instrumental to the widespread adoption of DR.

Aggregators are entities that act as large energy buyers and load reducers on behalf of a number of consumers. Regarding DR, aggregators typically bid wholesale load reductions with the utilities while negotiating and compensating its customers for the reductions. To the utility, an aggregator is seen as a large consumer. To the consumer, the aggregator is seen, in a sense, as the utility. In terms of ADR, the utility can send the reduction request to the aggregator, which, in turn will, forward it the consumers. The added value of aggregators lies in their knowledge of the specificities of the consumers on behalf of which they operate. In principle aggregator could also act as a capacity buffer performing load-balancing and some enrolling consumers in greater reductions to offset peak consumption of a set of consumers to uphold the overall reduction agreement.

It should come without surprise that a large building could be connected to an aggregator and itself act as a lower level aggregator for its occupants. Taking this concept further, one could envision a setting where each company in the building could act as an aggregator for each department. Upon receiving energy pricing information or load reduction calls, the




Figure 1. Overview of the SmartLink system illustrating a generic setup of nodes whereby buildings can be directly linked to the utility, connected to an aggregator, or may, in some cases act as an aggregator for its own occupants.

company’s aggregator would forward the appropriate reduction calls to each department that would, in turn, take the appropriate actions. At the lowest levels, the actions may translate into informing occupants, actuating automatically on devices or even lead to adjustments in business activities. The actuation on devices could primarily use the Building Automation System (BAS) or other specific system that controls electrical loads. This behavior should ideally take into account updated information regarding the current buildings’ state and consumer requirements.

To implement the above mentioned scenario we propose a hierarchical model for DR, which we designate as *SmartLink*. Our model is best understood as a tree of aggregator nodes, with compatible interfaces, linked to one another (see Figure 1). The top node corresponds to the interface with the utility, the inner nodes correspond to aggregators that encapsulate load-reduction strategies optimized to explore some kind of domain or context knowledge and, finally, the terminal nodes are the ones that interface with specific equipment to execute load reductions. Ideally, these actions should take into account local building, equipment and user constraints to provide an optimal trade-off between two conflicting requirements: user comfort and energy prices [10].

Our text is organized as follows: The next section overviews the concepts related to ADR systems. In Section III we present the SmartLink model and, in section IV, we discuss various application scenarios. Finally, in Section V we close by presenting the conclusions of this work.

II. AUTOMATED DEMAND RESPONSE

Currently available solutions for ADR are still very ad-hoc in nature, i.e., they are highly dependent on the building and on its particular equipment, and the type of requests sent by the grid. The response of the system and most idiosyncrasies of the automation hardware being controlled are actually hard-coded in the application software. The net result are solutions that are expensive to adapt to new buildings and lack the agility to evolve to accommodate new requirements.

Smart thermostats that relax the temperature setpoints at higher energy prices [11], [12] or appliance gateways that turnoff appliances such as washing machines or freezers are examples of these ad-hoc technologies. These approaches are

limited since the range of devices whose control can be handed over in this way is small. Moreover, these approaches are not adaptable since the load is controlled directly and there is little room to account for the user preferences or for changes in the *context* such as ambient temperature, luminosity, or the task being performed. In contrast, we propose a layered approach to DR whereby each node is responsible for determining the ideal actuation to fulfill the reduction requests.

A noteworthy remark is that the uncertainty regarding the true future needs may cause many customers to refrain from committing to certain ADR agreements. Second, although in some periods some customers may not tolerate the shutdown of one type of load but may well accept the shutdown of another. Third, there are cases where full shutdown may not be acceptable but some pattern of duty cycling or intensity control could be used to achieve the same overall consumption reduction. Whichever the case, we claim that an automated system should be able to decide on behalf of the user. Still, in other situations a full reduction may not be attainable but both the consumer and the utility could benefit from a partial commitment. Also, it is conceivable that, for some consumers, bigger reductions can be attained which compensate for smaller reductions from other consumers. To the best of our knowledge, this is not yet possible with the current technology.

Pilot studies like the EnergyPlus tests at the NY Times building indicate that ADR is technically feasible with existing BAS technology [10], [13]. The main sources of electrical load are, typically, HVAC and Illumination, which are commanded by the BAS. There is a large potential for using the building thermal mass for load shifting, using techniques such as pre-cooling by circulating cool air from outside within the rooms during the night or early morning at off-peak hours to benefit from a lower tariff. Field studies have shown that potentially there is a tremendous opportunity for reductions in on-peak energy. However, peak demand and the savings potential are very sensitive to the utility rates, building characteristics, weather conditions, and occupancy schedule [15]. Load shedding can also be applied for illumination systems through continuous dimming, for example, or alternating on/off in rows on circulation areas or open space areas. Yet, the creation of systems that connect ADR to smart buildings that successfully explores these techniques is still an open problem [14].

Another important aspect of an ADR system is the inexistence of commonly accepted and standardized interfaces that abstract the roles of *DR provider* and *DR client*. In particular, the utility should comply with the interface of a DR provider and the aggregator should comply with the interfaces of both DR provider and client. The systems that execute the actual reduction actions should comply with the interface of a DR client. However, even standardization initiatives such as OpenADR [16] define a specific role for the DR aggregators, which, in practice is neither compatible with a DR provider or client. For simplicity of presentation we will henceforth refer to *provider nodes* as being nodes that implement a DR provider interface and possibly aggregate multiple *client nodes*, which implement a DR client interface.

III. THE SMART LINK SYSTEM

SmartLink aims to be a generic highly customizable solution for fully automated DR offering a uniform interface for DR providers and client nodes. By highly customizable we mean a system where distinct types of aggregators may provide their specific implementation as long as they respect the interface. Conceptually, each node may encapsulate a consumption optimization strategy that takes advantage of the requirements of different businesses, of the specificities of buildings and of occupants' preferences. The definition of a standard and unique interface for DR provider and DR client opens a new array of possibilities. Since all nodes implement the same interface they can be easily chained, resulting in a hierarchy of aggregators that enroll and coordinate other nodes. This approach simplifies the development of large-scale multi-level ADR systems.

The hierarchical model that we propose promotes the separation of concerns, abstracting away the details of how consumption reductions are achieved at lower levels. Provider nodes may forward tariff information and load reduction requests down to its clients. In turn, clients may send notifications and information back to their corresponding provider. This information may be instrumental for monitoring and to estimate future consumption needs. Nodes may have different degrees of sophistication. Some may work in a fully automated manner, while other may rely on human experts to trigger their actions. Our model is abstract in this respect. The top represents the interface with the utility (i.e., the SG). Nodes at higher levels in the hierarchy may have more intelligence with respect to negotiation as they are expected to be more capable of taking advantage of evolving tariff plans that will be created by the utilities in the future. The terminal nodes are expected to act only as DR clients and will interact with the existing equipment and take into account actuation constraints as well as business and user comfort constraints.

Below we describe the services of an ADR system in terms of our hierarchical framework of provider and client nodes. Thereafter, we propose a communication protocol to link the nodes and specify the messages exchanged between nodes that implement the ADR services. We close this section with a discussion regarding the implementation issues of the different types of nodes.

A. Basic ADR Services

In our hierarchical ADR proposal we define a set of fundamental services that must be offered, which we detail in the following.

1) *Load Reduction Request*: The primary mechanism of ADR is informing clients of load reduction requests. Reduction requests can be mandatory or optional. A mandatory request often entails a penalty to the client if not executed. In optional requests, clients may decide to participate based on trade-off analysis made between pricing information and business impact estimations, which can include the impact on user's comfort. This information is expected to change, leading the node to review its decisions. A reduction request may arrive at a pre-determined time or

asynchronously, possibly to accommodate an unforeseen load reduction need of the provider. A reduction request is also characterized by its starting time into the future, which can vary in the order of minutes to days. Ideally, greater reduction requests should arrive with sufficient time ahead to allow for an adequate planning.

2) *Load Reduction Execution*: Client nodes must have a mechanism to carry out load reductions at specific time intervals. The implementation of this mechanism is abstracted in our model and can be of two types. In the case of an aggregator node, the typical behavior consists of delegating the reduction further down to its clients, possibly coordinating them to perform load balancing. An aggregator may send distinct reduction requests to its clients according to some internal pre-defined balancing policy. At the end of the chain, terminal nodes will typically command loads considering user preferences and constraints.

3) *Load Reduction Verification*: Each provider node must have a way of ascertaining whether the clients as agreed have performed the agreed reductions. In principle this is achieved through metering and serves two main purposes: (i) validate the compensation assigned to client nodes for their participation and (ii) detecting exceptional situations where a client node cannot fulfill the previously agreed reduction. Detecting exception situations is important, as the node may need to perform adjustments to its balancing policy requesting other nodes to compensate for the unforeseen consumption. In this way the node itself may be able to avoid violating its own accepted reduction request.

B. Advanced ADR Services

Certain ADR functions require more sophisticated services to be implemented.

1) *Future Consumption Estimation*: Estimated future consumption information given by clients is useful to fine-tune the consumption estimates of the provider making it more effective. A node may use this information to manage and plan future consumption levels and, depending on its implementation, it may schedule generation, send load reduction stand-by requests down to clients or perform load shift or load shedding actions. A provider may aggregate information from multiple clients and then make decisions based on it. In our model, a node supplies its estimations to its provider, thus propagating information upwards in the hierarchy chain, enabling upper nodes to achieve global estimations. Alternatively, decisions can also be based on other factors. A building provider node may consider, for example, statistical data and process short-term events regarding building usage, consider scheduled meetings, detect users' arrival or departure, detect environmental or weather changes. In this way, more educated estimates can be produced.

2) *Reduction Potential Estimation*: From a management point of view it is very important that provider nodes know how much power reduction they can engage. The reduction

potential can be estimated by asking client nodes for their baseline consumption, i.e., the consumption that would be in effect if no reduction request is being executed, and then asking for the power consumption estimates that they foresee to achieve for a certain reduction effort.

3) *Bidding*: Demand Bidding can be implemented by having provider nodes announcing energy prices to clients, where the energy price can vary between a default value and a minimum value. Each client node can respond with a proposal in which specifies the energy cost it will accept to execute the requested reduction. The provider will then notify or refuse the proposal. This bidding process can be executed down the hierarchy chain. We anticipate that the lower level nodes are less likely to have bidding capabilities.

4) *Structured Interruption of a Reduction*: Due to unforeseen or emergency events, a node may not be able to uphold a reduction or a tariff it may have previously committed to. Since monitoring can be deferred in time, a mechanism must exist for the client to inform the provider that the agreed reduction will no longer be met, thus allowing the provider to handle the interruption and take useful action in sensible time. An appropriate action to accommodate the increase in consumption associated with the interruption notification is to trigger a new load balancing round between client nodes. If the provider node determines that it cannot accommodate interruption, perhaps because too many other client nodes have also requested an interruption as well, it will need to send an interruption notification to its own provider, propagating the interruption upwards in the chain. Interruptions can be accepted by the utility or by the aggregator, with a penalty to the client node. Another way for a node to deal with interruptions is to manage a consumption reduction reserve capacity. However, maintaining a high level of reserve capacity limits the efficiency of the ADR system as a whole.

C. Node Protocol Overview

In this section we propose the messages that the nodes will use to interact with each other.

Smartlink message exchange is organized according to three phases. The first phase is the *planning phase* where nodes exchange information regarding expected energy prices and consumption estimations. The second phase is the *commitment phase* where client nodes commit to a given tariff or reduction request, enabling provider nodes to compute better estimates of their reduction potential. The last phase is the *execution phase*, where client nodes are asked to actually perform a reduction and are monitored by the provider nodes.

The planning phase provides for tariff announcements as well as for the estimation of future consumption and for the estimation of reduction potential. Tariff announcements consist of tuples sent by providers. A tariff tuple takes the form $\langle t_i, t_f, p, c, c_{min} \rangle$ where $t_i < t_f$ are the initial and final timestamps that define the tariff effect interval, p is the absolute power limit in KWh, c is the energy baseline cost and c_{min} is the lower cost bound. Whenever bidding is not available c_{min} matches c . An

announcement can override another, in which case their respective initial and final timestamps are the same. Multiple announcements can be issued to the same combinations of t_i and t_f , thus offering different rates for distinct power consumption ranges. Any node upon receiving these announcements can prepare to take advantage of a given lower tariff by undertaking the appropriate actions such as preparing for load reduction, scheduling on-site generation, or in the case of aggregator nodes, passing down the tariff to client nodes.

We assume that different arrival rates of tariffs may exist. Higher rates will be needed to support greater variability in dynamic pricing. Notifications can be classified in short-term vs. long-term depending on the interval between the current time and t_i .

The notification of future consumption and reduction potential is achieved by the clients informing providers about their expected consumption on a given time-frame by issuing tuples in the form $\langle t_i, t_f, p, o \rangle$. These tuples inform that a total power consumption of p , in the specified time interval, is expected to have an impact o . The impact is a global discrete value shared and understood by all nodes. The value *NONE* is used to inform about unconstrained or regular operation, *MIN* means a saving scenario with marginal disruption of business or occupant comfort, *MAX* is a scenario of maximum achievable savings, possibly implying business changes or occupant discomfort. Finally, the value *CRITICAL* is used for emergency or catastrophe situations that require reducing power to the absolute minimum to avoid a global shutdown of the power network.

In the commitment phase, clients inform their providers of their desire to commit to a certain previously announced tariff. Clients send an acceptance tuple regarding a tariff in the form $\langle t_i, t_f, p, c_{new} \rangle$. To signal commitment (i.e., participation) the price c_{new} can be equal to price c of a previously announced tariff. However, if $c_{new} < c$ then c_{new} is a new price proposal, i.e. a bid, which is being proposed by the client. Bids have to be acknowledged by providers by sending clients an acknowledge tuple in the form $\langle t_i, t_f, p, c_{new} \rangle$.

In the execution phase providers send clients execution requests in the form $\langle t_i, t_f, p, m \rangle$ asking them to comply to the maximum power p in the specified time frame. The parameter m is a flag indicating if the execution is optional or mandatory. A provider may issue a new tuple $\langle t_i, t_f, p' \rangle$ for the same time frame, increasing or decreasing the reduction effort, respectively, if $p' > p$ or $p' < p$. If the client needs to opt-out it may inform its provider by issuing a tuple in the form $\langle t_i, t_f, t_{out}, p_{new} \rangle$ where $t_i \leq t_{out} < t_f$. The timestamp t_{out} refers to the instant after which the client is expected start consuming at least p_{new} .

The support for monitoring is achieved by clients issuing tuples in the form $\langle t_i, t_f, p \rangle$ where p is the actual power consumed in the interval (t_i, t_f) .

D. Implementation issues

The implementation of each type of node requires a similar set of issues to be addressed, which we now briefly overview.

1) *Organizational policies*: Some aggregator nodes can be responsible mainly for propagating ADR information to client

nodes accordingly to policies defined by the building owner, the facility manager, or other responsible persons. To perform this task, nodes need to be configured adequately and have access to information that allow them to implement specific reduction demand balancing policies and issue reduction requests for the lower level nodes. These policies can be written in a tailored rule-based language and interpreted by nodes that require this functionality. Another important function that high-level nodes may require is automated bidding capabilities, which needs also to be adequately defined and configured.

2) *Building description*: Certain nodes node may be associated with parts of a building. It can be the complete building, grouping of different areas such as a corridor, a whole floor or the rooms belonging to a given tenant or even a part of a room. Certain nodes can undertake energy optimization decisions based on data about the building itself like room orientation or materials applied.

3) *User preferences and constraints*: Nodes that aim at to accommodating user preferences and constrains must have access to user profile information that reflects the users' preferences for the space according to the type of task that is to be performed therein. For example, the node must be capable of determining what is more relevant to the user (good light, precise temperature, etc), the acceptable adjustment margins (light can vary between 50 and 90%; temperature set point can vary 1.5C), and what restrictions exist (some equipment cannot be turned-off; or lighting cannot be reduce during certain tasks).

4) *Interfacing with equipment*: An automated system to assist in demand response will have to interface not only with the system controlling the loads but also with the operational system supporting the business processes of the occupant of a building. Note that the terminal nodes (lowest level nodes) are nodes that interface with specific equipment and issue direct commands to loads. It will have a generic interface that maps reduction requests into the actions support by devices (linear actuation vs. discrete actuation). These generic actions will then be forwarded to specific drivers associated with each device and translate into specific commands of the device's technology. These nodes will have to maintain a list of equipment (loads) and their detailed characteristics that are associated with it.

IV. APPLICATIONS

This section discusses several application examples, which aim at illustrating our proposal potential.

The simplest scenario we can think of consists of performing direct load control into a home with a heating and cooling system engaged in a demand response program with its corresponding utility (see Figure 2). In this example there will be just two nodes (a client node) that interfaces with the utility (the provider node) and receives the reduction demand requests. Accordingly with user preferences and to other defined policies (which may include pre-contracted arrangements), this node will command the acclimatization




Figure 2. Basic scenario implementing a direct load control scheme.

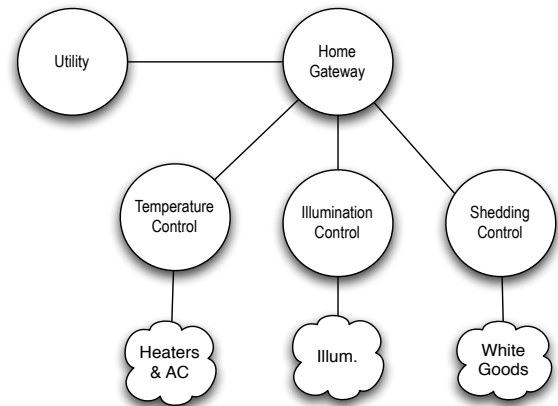


Figure 3. A smart home scenario with a home gateway acting as an aggregator.

equipment and fulfill the utility requests. This very simple application implementing the DR client interface mimics a “smart thermostat”. However, our approach can offer expandability capabilities and can implement a richer service than just direct control.

Another example is considering a system with a two-level hierarchy of nodes, depicted in Figure 3. The top aggregator node interfaces with the utility and balances the reductions demands into three terminal (client) nodes that command, respectively, an acclimatization system, the home illumination and a set of white goods apparatus. This scenario supports a richer set of options, where user preferences and restrictions can be accommodated accordingly with different usage patterns of the home in a distributed fashion by each of the terminal nodes. For example, the temperature and lighting control can consider space usage, day time period and user activity, to balance reductions in a way that minimizes user impact. The terminal node associated with white goods control may, among other functions, manage load shedding through time shifting.

The next example, shown in Figure 4, illustrates a large building, which may have different tenants or lodge just one big company with different departments. We suggest a three level hierarchy of nodes. The top-level aggregator node interfaces with the utility/smart grid and will be responsible for the whole building behavior and response to reduction requests. In this scenario, the overall building response will be delegated

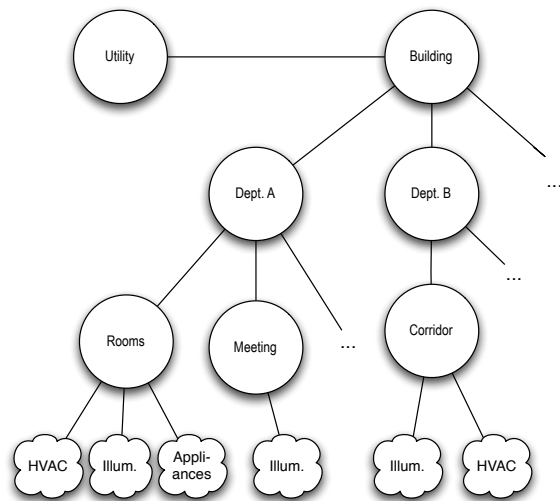


Figure 4. Basic scenario implementing a direct load control scheme.

to different departments, managed by their own aggregator nodes, accordingly with some pre-defined policies that will guide balancing between the next level nodes. Each one of these nodes will then map into to specific building functional areas, which may be, for example, individual rooms, open space areas, corridors, and meeting or leisure areas. The nodes responsible for managing a building may take into consideration the user preferences and possible restrictions and interface with terminal nodes that will control specific loads.

As a last example, we consider multiple aggregator entities that isolate the utility/smart grid from the final consumers. We envision a multi-level hierarchy where the utility/smart grid interfaces with a town aggregator entity which, in turn, interfaces with various other aggregation entities associated, for example, with neighborhoods, which in turn will interface with individual buildings. Of course each building can have its own aggregator node as described in the previous example.

V. CONCLUSIONS

We believe that the current models of ADR are too limiting for both the utility and the customers and, in practice it is hampering the potential of ADR. In this paper we propose the concept of decoupling the reduction requests sent by the grid from their implementation through a hierarchical ADR model—SmartLink—that is based on a unified interface for demand response providers and clients.

Each node hides details from its provider encapsulating consumption reduction strategies that take advantage of the particular context associated with each consumer. Our approach offers a way of structuring a global system in a hierarchical way making it more scalable and manageable. We

discussed how relevant ADR services could be implemented in our hierarchical model. Moreover, we proposed a protocol for nodes to interact with each other and illustrated possible application scenarios for our model.

However, creating such an ADR system presents significant technical challenges and requires a multidisciplinary cooperation of the Software Engineering, Automation and Smart Grid communities. Only then we can expect a broad acceptance and dissemination of this type of technology.

REFERENCES

- [1] A. Ipakchi and F. Albuyeh, "Grid of the future," *Power and Energy Magazine*, IEEE, vol. 7, no. 2, pp. 52–62, Mar.-Apr. 2009.
- [2] A. S. Massoud and B. Wollenberg, "Toward a Smart Grid: Power delivery for the 21st century," *Power and Energy Magazine*, IEEE, vol. 3, no. 5, pp. 34–41, Sep.-Oct 2005.
- [3] M. Albadi and E. El-Saadany, "Demand response in electricity markets: An overview," in *IEEE Power Engineering Society General Meeting*, 2007, Jun. 2007, pp. 1–5.
- [4] A. K. David and Y. Z. Li, "Consumer rationality assumptions in the real time pricing of electricity," *IEEE Proceedings of Generation, Transmission and Distribution*, vol. 139, no. 4, pp. 315–322, 1992.
- [5] A. Faruqui, D. Harris, and R. Hledik, "Unlocking the 53 Billion euro Savings from Smart Meters in the EU," Report by The Brattle Group, Inc., Oct. 2009.
- [6] M. A. Piette, O. Sezgen, D. Watson, C. S. N. Motegi, and L. ten Hope, "Development and Evaluation of Fully Automated Demand Response in Large Facilities," *California Energy Commission*, Tech. Rep., 2005.
- [7] S. Braithwaite, "Residential TOU Price Response in the Presence of Interactive Communication Equipment," in *Pricing in Competitive Electricity Markets*, A. Faruqui and K. Eakin, Eds. Kluwer Academic Press.
- [8] N. Hopper, C. Goldman, and B. Neenan, "Demand Response from Day-Ahead Hourly Pricing for Large Customers," *The Electricity Journal*, vol. 19, no. 3, pp. 52–63, 2006.
- [9] A. Faruqui and S. Sergici, "Household response to dynamic pricing of electricity – A survey of the empirical evidence," *SSRN eLibrary Working Paper Series*, 2010.
- [10] S. Kiliccote and M. A. Piette, "Advanced Control Technologies And Strategies Linking Demand Response And Energy Efficiency," in *Proc of Int'l Conf. for Evaluation of Building Operations (ICEBO '05)*, 2005.
- [11] M. Martinez and F. Russell, "Smart thermostats getting results at SCE," *Utility Automation & Engineering*, Jul 2007.
- [12] T. Woody, "From Smart Meters to Smart Thermostats," *NY Times*, November 2009.
- [13] M. A. Piette, D. Watson, N. Motegi, and P. X. S. Kiliccote, "Automated Critical Peak Pricing Field Tests: Program Description and Results," *Berkeley National Laboratory*, Tech. Rep., 2006.
- [14] H. Chen, P. Chou, S. Duri, H. Lei, and J. Reason, "The Design and Implementation of a Smart Building Control System," in *IEEE International Conference on e-Business Engineering (ICEBE '09)*, Oct. 2009, pp. 255–262.
- [15] J. E. Braun, "Load Control Using Building Thermal Mass," *Journal of Solar Energy Engineering*, vol. 125, no. 3, pp. 292–301, 2003.
- [16] M. A. Piette, G. Ghatikar, S. Kiliccote, E. Koch, D. Hennage, P. Palensky, and C. McParland, "Open Automated Demand Response Communications Specification (Version 1.0)," *California Energy Commission*, Tech. Rep., 2009.