Social Control of Power System Demand Based on Local Collaborative Preferences

Michael N. Huhns

University of South Carolina, Dept. of Computer Science and Engineering, Columbia, SC USA
huhns@sc.edu

Abstract. This paper describes a computational approach to energy use that assigns importance to human psychology and social interactions. Specifically, this paper describes our investigations into computational mechanisms that encourage prosocial behavior on the part of consumers. Examples of prosocial behavior in the context of electrical energy use are reducing average aggregate consumption and peak total consumption. We consider an approach that combines minority games and cake-cutting that includes elements of human decision-making in situations that are hybrids of competitive and cooperative settings. For example, people may be motivated to reduce their consumption if that were posed as a competition wherein they would win a game, possibly by collaborating with their neighbors. And, people may be motivated to behave in a prosocial manner if selfish behaviors were shunned in their social group. Previous approaches disregard such dynamics from technical studies, relegating them to psychological analyses; yet the interrelationship of the human and the technical aspects is crucial in a complex sociotechnical system such as the power grid.

Keywords: Multiagent systems, electric power, demand-side control, social computing.

1 Introduction

There are many facets to the world-wide electric power problem, concerning how electric power can be generated in an environmentally sound way, how it can be stored and distributed efficiently, and how it can be used wisely. Although energy resources can be viewed strategically as an advantage for geopolitical entities that own the resources, we prefer to view them broadly as societal resources to be shared among the members of a society. The *control* of energy resources is *not* societal, however: it is centralized at the energy provider, where preferences of the members of the society are generally not considered.

We are investigating the modulation of electric power demand via socially intelligent computing. We seek to develop efficient consensus and incentive-based computational mechanisms for decentralized control of demand that respects system-wide objectives and individual preferences. Our mechanisms will influence consumer decisions regarding local energy usage, generation, and storage, as well as overall energy supply and demand, according to local consumption preferences and global supply

adfa, p. 1, 2011.

objectives of grid operators. The societal benefits are lowered peak demand, improved operating efficiency, and lowered capital expenses.

1.1 Current Situation

In general terms, our problem involves the allocation of electric power (treated as a scarce societal resource) among independent consumers (households and small businesses). Recent approaches collected under the term "smart grid" enable consumer devices to be controlled by electric power utilities. The objective is to shed demand when it exceeds supply. For example, household air conditioners can be turned on or off easily from a central controller. However, deciding *whether* and *when* to turn them on or off based on consumer preferences is nontrivial. The smart grid is smart only from the viewpoint of the electric power utilities. Because consumers typically want their preferences to remain private, centralized approaches that allocate resources by fiat are not acceptable. How can consumers with diverse preferences make local decisions about the allocation and management of electric power that are globally effective? The problem is exacerbated by large consumer communities and frequently changing preferences.

Two forms of demand-side management are being used to solve energy resource allocation problems. In one, a central control form, a utility enters into agreements with customers, for a rebate incentive, under which the utility can directly control appliances, usually for load shedding when needed [1]. Central control does not address customer comfort and exception requests. In the other form, home management systems monitor and manage appliances. Some utilities are considering providing real-time pricing signals to improve the effectiveness of such systems. Home management systems suffer from customer reluctance to participate and lack of clear benefits [2]. Utilities have begun to realize that pure technical or economic approaches are not effective, so they are investigating alternatives to better engage their customers [3, 4].

Here is an example of the problem we are trying to address. Charging an electric vehicle is equivalent to approximately four houses using all of their appliances. The transformers in a neighborhood (the ones you might typically see on a pole) are sized for approximately ten houses. If 3 or 4 people in a neighborhood buy an electric vehicle and try to recharge it at the same time, the transformer will fail. To prevent this, the power company could double or quadruple the capacity of their lines and transformers, which would be very expensive, or the power company could take control of when people can recharge their vehicle, OR the neighbors could cooperate with each other in staggering when they recharge. We believe that the last is the best solution, but it requires the neighbors to be cooperative and possibly altruistic, and it must be done with local consumer cognizance of the global context.

1.2 Investigation Framework and Research Hypotheses

Our investigations are being conducted in a framework of realistic premises designed to make this large problem manageable. The premises are

Premise 1. Current pricing incentives are insufficient, because they are based on a history of past aggregate behavior and have little predictive value.

Premise 2. The community of consumers exhibits rich social relationships and energy usage dependencies that can be handled better through peer-to-peer interactions rather than through centralized control.

With these as a basis, the key is fostering peer-to-peer interactions among consumers to guide their individual control decisions and, by aggregating the decisions, produce effective system-level control. Individual demands are coordinated to reduce peak demand, flatten overall demand, and yield a power factor closer to 1.0. We believe that two levels of peer-to-peer interactions will be needed. At the macro-level, interactions create consensus on the overall goals and trade-offs, producing the equivalent of supply-and-demand curves. At the micro-level, interactions cause individual control decisions to be as dissimilar as possible, so as to spread demand as uniformly as possible.[1] To investigate this foundation for a demand-side approach, we have formulated the following hypotheses:

Hypothesis 1: Participation. A sufficient number of people in a society can be motivated to participate either directly or indirectly via their intelligent software agents in the management of an essential and limited resource (electric power).

Subhypothesis 1.1: Influence. Consumers' decisions can be influenced to promote prosocial behavior, if such behavior does not detract from their personal preferences.

Subhypothesis 1.2: Privacy. Consumers will share some private information (indirectly via their agents) so as to cooperate in promoting prosocial behavior.

Subhypothesis 1.3: Cooperation. Consumers are more amenable to promoting prosocial behavior if they can cooperate with known parties, not with anonymous strangers. Consumers who cooperate will achieve better outcomes.

Subhypothesis 1.4: Competition. A game environment offering competition among consumer groups can motivate consumers to exhibit prosocial behavior.

Subhypothesis 1.5: Trust. Consumers will trust software agents to represent their interests in negotiating for resources.

Hypothesis 2: Stability. A system of interacting agents cooperating and competing for resources on behalf of a community of users will produce a controllable, stable, and prosocial allocation of resources.

The scientific results will be improved understanding of how the macro-level and micro-level aspects of control come together and how users remain in control while engaging in socially desired behaviors. There will be three interdependent types of macro-level and micro-level interaction among providers, consumers, and the software agents representing individual interests (see Figure):

1. Expressing preferences: software agents interact with consumers to acquire their preferences and provide incentive-based feedback to influence their behavior.

---

[1] We also recognize that at times it is important to develop "herd behavior." For example, if power is largely solar, then it is preferable to use that energy as it is being produced. Or, if industry needs large power during working hours, then it is desirable to push all residential demand into the nighttime so that the net consumption becomes flatter.

2. Reaching consensus: macro-level interactions among agents to optimize their distributed demand decisions based on computational collective intelligence and consensus-based optimization, resulting in supply-demand curves based on local preferences and system objectives.
3. Achieving objectives: micro-level interactions between groups of consumer agents and resource provider agents to minimize impact on resources (i.e., reduce peak demand) based on field theory from particle physics, cake-cutting algorithms from studies on negotiation among multiple agents, and incentive-based optimization mechanisms.

The problem of allocating shared resources matches naturally with socially intelligent computing—the intersection of social behavior and distributed computational systems—and multiagent systems. Multiagent systems can apply social computing to investigate the technical problem of how to allocate, distribute, and govern scarce societal resources in a sustainable manner across a sufficiently coherent community of users, each potentially having different preferences for the resources and when to consume them. This is difficult, because communities are large and preferences will change frequently and, from a centralized viewpoint, unpredictably. The sharing should accommodate member preferences, yet provide fair and envy-free incentives to those whose preferences most promote sustainability. Because preferences are mostly personal and private, centralized approaches that allocate such resources by fiat are usually not acceptable in a free society. The key to this is fostering peer-to-peer interactions among the participants so as to accommodate both the personal and the interpersonal dimensions of decision making by rational, social participants.

The two complementary strategies we have investigated for the interactions are based on: (1) control systems and (2) a negotiation approach that combines minority games, particle physics, and cake-cutting algorithms. Multiagent negotiation is one of the decision-making mechanisms that can be used to provide for the allocation of resources. The results of such negotiations, from the viewpoint of the consumers, should be fair and envy-free, which motivates the investigation of "cake-cutting" approaches.

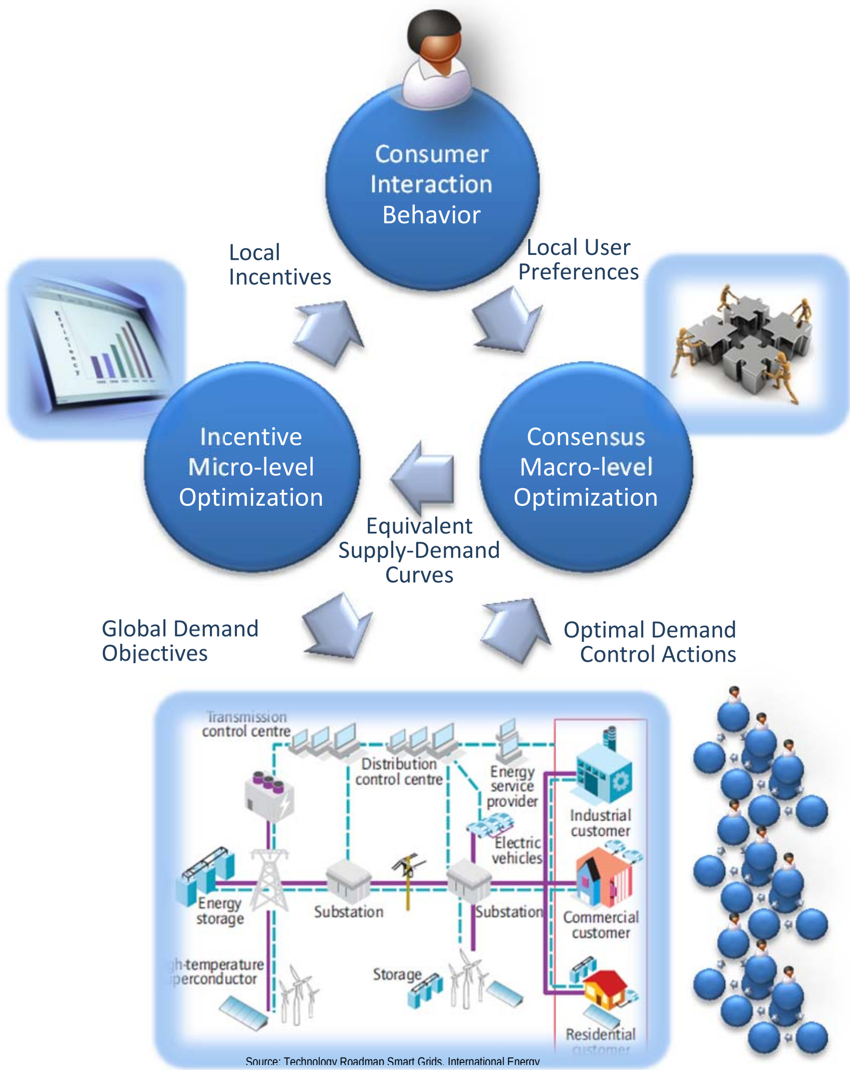

2 Background and Significance

Crowdsourcing [5] has drawn a lot of interest lately. Crowdsourcing involves (usually implicit) collaboration between users to solve a problem. However, crowdsourcing approaches are fundamentally limited to solving centrally allocated problems where the mode or median or individual solutions converge to the ideal solution. We refer to such central tendencies loosely as the majority. In majority problems, a statistical aggregation of individual solutions proves effective. To follow Galton's example from 1907 of the wisdom of crowds, if 800 people estimate the weight of an ox by just looking at it, their individual estimates may vary a lot, but in a Gaussian manner: their majority estimate (the median) could be close to the actual weight of the ox. So much so that when we have no other means to determine the weight of an ox, we might rely solely on the majority estimate, which is what crowdsourcing pursues. Such solutions can be promoted by giving users an incentive to be nearest to the majority. Notice that if we gave users an incentive to be far from the majority, the result would be meaningless.

In the case of resource usage, however, the participants' interests are not well aligned with the majority. If increased peak demand causes the price to go up, consumers are better off spreading their individual loads to lower the peak and, thus, lower the price. In such settings we are not seeking a majority view of the "right" time to consume energy, but to influence consumers to distribute their consumption. The consumer in effect has an incentive to be in a minority. Minority settings in general are highly volatile. Are there social mechanisms that can motivate behavior to produce effective solutions in minority settings?

The *minority game* [7][9] is defined as a game with a large number of players, N, with each player making a choice between two alternatives at each round of the game. After all players have made their choice, the players that are in the minority each win one point. This is relevant for electric power distribution, because the preferred solution is for consumers to request power at different times.

A variant of the minority game is the Kolkata Paise Restaurant Problem [8] where the number of choices (n) as well as the number of players (N) are relatively large. It is a repetitive game where information regarding the history of choices made by different players is available to everyone. Assuming that $n = N$, a player ϵ N wins a point by making a choice ϵ n made by no other player. If a choice is made by more than one player, one is randomly selected to earn a point. Hence, while each player gains a point for making a unique choice, the resource utilization is maximized when each choice is made by at least one player.

2.1 Power Systems and Control Theory

Current demand-side management approaches fall into two main categories. First, in central control, the utility enters into agreements with customers, for a rebate incentive, under which the utility can directly control appliances, usually for load shedding when needed [1]. Central control does not address customer comfort and exception requests. Thus, customers are reluctant to participate and only a few do. Second, home management systems monitor and manage appliances (e.g., by turning them on and off, or adjusting temperature settings). The consumer is expected to play a major role in (paying for) installing and configuring such systems. Utilities can provide realtime pricing signals to improve the effectiveness of such systems. Walker and Meier [2] observe that home management systems suffer from customer reluctance to participate, and lack of clear benefits. They also observe that some kind of automation is essential: as they found, in settings where consumers sought to control their air conditioner usage manually, they would turn on their respective air conditioners precisely at peak times thus exacerbating peak demand.

As support for the significance of our proposed approach, Berst [3] points out that pure technical or economic approaches are not proving effective and utilities are investigating alternatives to better enlist their customers' support. Similarly, the Federal Energy Regulatory Commission [12] acknowledges the challenge of communicating the importance of demand response and engaging consumers effectively.

In recent influential works, Sean Meyn [4] has articulated well some of the challenges of relying purely on pricing mechanisms for system control. At the macro-level, such approaches have led to well-known problems though they have demonstrated that consumers can change their demand in response to severe price signals. However, this doesn't mean that the resulting allocations are equitable or that consumption is smoothed in the process. Mathieu et al. [14] study different types of industrial and commercial consumers and observe challenges in prediction, specifically, that variation may often be dominated by model error rather than due to explicit response. Shao et al. [15][16] study residential load profiles, including the charging of electrical vehicles, which creates heavy loads. Shao et al. are concerned with capturing consumer priorities regarding various appliances and being able to control them as a way to shape the overall load.

Japan's Digital Grid Consortium envisions large-scale energy grids that can handle power the way the Internet handles data, using routers and service providers to efficiently direct the flow of electricity [17]. The consortium seeks to develop technology that can track units of energy across an entire grid, tagging them with their source and destination similar to the way Internet packets are handled. The consortium plans for inputs to include existing power plants, solar facilities, and other alternative sources. The grid will include local power storage systems, such as large-scale batteries in homes. The units of energy will be managed by service providers, tracing and charging for them like a currency exchange. The energy "messages" are intended for *supply-side management*, but could be adapted to serve *demand-side management*.

Because power systems are inherently distributed, agent-based approaches apply naturally therein to support local control. They contrast with extant approaches, which develop centralized solutions, placing all the intelligence in central controllers. Baran and El-Markabi [18] show how to characterize a multiagent protocol that facilitates control in the presence of local sensitivities as long as appropriate communication assumptions are met. Hernández et al. [19] study the modeling of power sources in smart grids. Pipattanasomporn et al. [20] apply multiagent systems from the utility standpoint. They show how their approach can isolate a local system from the grid adaptively as needed. Pipattanasomporn et al. [21] motivate a home power network architecture, which accords with our conception. Their proposed home management system corresponds to an agent that controls local loads on behalf of a consumer and responds to signals from the grid.

2.2 Multiagent Negotiation for Multiplayer Resource Allocation

An important feature of multiagent systems is that the agents can behave autonomously considering the interest of the people they represent. Fairness and envy-freedom are criteria used to judge the effectiveness of allocation procedures. Assume the resource being allocated is measurable. An allocation procedure is called fair if it distributes a resource among *n* agents such that every agent values its portion as exactly *1/n* of the total value of the resource. An allocation procedure is called envy-free if every agent values its portion at least as much as the portions allocated to other agents. Thus, envy-freedom is stricter than fairness. When a mediator is involved in resource allocation, an additional desirable criterion is that the mediator is unbiased. In addition, the procedure should be efficient in time and space complexity, strategy-proof, and constructive.

In open multiagent systems there is generally no global control, no globally consistent knowledge, and no globally shared goals or success criteria [22]. So the agents compete to maximize their own utilities. We assume each agent's utility function is private. A negotiation protocol should be immune to information hiding and lying by the agents. In addition, protocols can be evaluated on various criteria such as fairness, envy-freedom, equitability, and efficiency. Brams and Taylor [23] discuss various procedures for allocating resources. They show that it is generally difficult for any given procedure to fulfill more than two of the above mentioned criteria. These criteria are by no means exhaustive, but may be taken as an initial test of the allocation procedure that is being proposed. For example, other criteria that can used to evaluate protocols are: simplicity, computational complexity, and verifiability.

A protocol for negotiated resource allocation—the basis for the multiplayer game envisioned here—is said to be verifiable if the allocation of the resource is invariant to the bias of the mediator (game engine). Iyer and Huhns [10][11] address verifiability in a resource allocation procedure for one or two-dimensional resources, proving that if the agents follow a specified multiagent negotiation protocol, it is possible to have a fair and unbiased allocation of the resource. At the end of the negotiation, one of the agents volunteers to act as a mediator and executes the procedure. Based on the computation of agent preferences, there are two outcomes: the procedure (i) finds a solution and all agents get a fair deal; or (ii) fails to find a solution and all agents receive the conflict deal, i.e., no agent receives any part of the resource. The salient point is that the agents can detect if the mediator attempts to manipulate the results. Hence the results of this method are verifiable to any agent who wants to check them and the mediator need not be a trusted outsider. Importantly, the utility functions of the agents are not compared and therefore are unconstrained: all that matters is how the agents' preferences relate to one another.

3 Analysis

Let us consider one concrete scenario of how sustainable energy use can be treated as a societal problem. This scenario seeks to reduce peak demand but does not address reducing total demand. That is, we would like consumers to shift their individual demands in time so that the peak aggregate demand at any time is reduced. Doing so has benefits in yielding a more stable load and reducing the need for capital expenses.

The most traditional approach would be to determine the popular times of the day or week for demand and to set the price higher for such times, so as to encourage consumers to move away from such times. Such an approach works from historical data and lacks knowledge of and flexibility in addressing changing situations.

A more modern approach is to apply real-time pricing. However, real-time pricing is difficult from a practical standpoint because of characteristics of power systems that cannot match production to demand instantaneously. Further, real-time pricing is difficult for consumers to deal with, and often leads to chaotic outcomes [5].

We now describe the interactions between a consumer, a power supplier, and our proposed system, which could be thought of as a mapping to an energy service provider (ESP) [13]. Let's begin with a variant wherein the consumers act independently of each other.

As the Figure shows, a consumer assisted by an agent submits constraints on an expected future load profile. A local broker/manager considers all the submitted profiles and determines a nonbinding allocation for each consumer that reduces the peak demand and demand variations. The allocation is guaranteed to satisfy each consumer's stated constraints. A simple way to find such an allocation is to order the consumers randomly and, for each consumer in turn, allocate power usage timeslots to that consumer in a way that greedily minimizes the peak consumption. Each consumer may or may not act according to the allocation.

A consumer who follows the recommended consumption profile pays the average price for the current total demand in each time slot. A consumer who consumes power arbitrarily either by never participating in our approach or by participating but deviating from the recommendation pays the usual marginal rate.

The price for power increases with the instantaneous demand at the time of consumption. With some key assumptions, this scenario provides a way to address some important properties:

- **Prosociality.** The local broker/manager charges a higher price to ad hoc consumers than for plan-ahead consumers, which creates an overall incentive to reduce peak demand.
- **Individual rationality.** Those who submit a profile and follow the resulting allocation benefit by paying a smaller price for the power they draw. Thus participants pay a lower price for power in a given slot than someone who consumes the same amount of power in the same slot but without a prior submission. Thus consumers are motivated to participate in the brokering and management.
- **No coercion.** Those who submit a profile are free to ignore the suggested allocation. They pay the same price for that consumption as if they had never participated.
- **Budget balance.** When consumers as a group create more expensive demand on the power source, they pay more for the privilege.

However, this approach assumes the consumers have NO knowledge of the constraints (preferences) of other consumers or of the constraints of the power generation and distribution system. That is, a consumer might be willing to shift its need for power to a slightly different time interval if it would result in a major savings in cost, but has no way of discovering this. This approach forms an imperfect information game.

Our approach combines pricing with social mechanisms. Consumers join cooperatives, which we assume are small, such as neighborhood blocks. Each cooperative seeks to minimize its overall cost in terms of financial units or in terms of environmental impact. Thus the members of each cooperative, must negotiate with each other with respect to their individual preferences as such preferences are affected by important externalities such as the changing price of energy, changing weather, and social factors such as whether it is a holiday season.

We are investigating some key challenges that arise from our vision, such as power system models, social interaction models, design models of agents, user models, and economic models. In addition to formal models and simulations, we are using games to explore how consumers interact in different circumstances and how we may effectively promote prosocial behavior. The interactions among power consumers might take the form of

1. Auctions, with the following features:
   - Individuals base their bids on their own preferences
   - Individuals do not reveal their preferences
   - Individuals could maintain and use a history of interactions. Based on this, individuals could learn the strategies of others, although the auctions might be designed to reduce or eliminate the need for this
   - The auctions do not allow any future considerations
2. Round-robin power scheduling, where individuals take turns having first preference for power use, in an endless cycle.
3. Direct negotiating among consumers, involving promises / commitments for future use, and which might be multiparty.

Particle physics provides both a metaphor and a mathematical basis for solving the resource allocation problem. Particle physics dictate that particles tend to occupy the most energetically favorable states, while certain other particles cannot occupy the same state together. This translates into an analogy of electric power resources that either any number of consumers can share or that only one consumer can have.

4 Research Agenda

The goal of the power company (maximize profit) is different than the consensus goals of its customers (minimize cost, maximize comfort, protect environment). Although the proposed project studies household electrical power consumption, its results could be applied to a broader class of societal resources, such as fresh water, thus promoting sustainability in such settings as well. Our approach applies social computing to sustainability problems. We treat consumers and providers as important participants and rely upon their mutual interactions—mediated by computational agents—as a basis for arriving at high-quality solutions. Each user delegates some authority to an agent, which then acts on the user's behalf. Traditional social computing approaches are limited to information problems where consensus is important. In contrast, our approach applies to allocation problems where the *dissimilarity* of the participants' decisions improves social welfare and helps capture each participant's local preferences. There are two main considerations:

1. Can a sufficient number of people in a society be motivated to participate either directly or indirectly via their intelligent software agents in the prosocial management of an essential and limited resource (electric power)?
2. Will a power distribution system managed from the edge by consumers be controllable and stable in a control system theory sense?

4.1 Uninvestigated Hypotheses

The following hypotheses are relevant and deserving of investigation, but this has not yet been done:

- Bottom-up preferences negotiated among users in a neighborhood are more secure than top-down control of appliances by power companies, as is envisioned for various "smart grids."
- Being cognizant of global warming and climate change, people will act altruistically towards their neighbors in allocating electric power resources.
- It remains to be shown that the grid will be more efficient and more fair if consumer preferences are considered.
- Because they have only local information and minimal global information, consumers have been shown to act suboptimally when the global grid is considered and not even in their own best interests locally when allowed to participate in decisions about the distribution and usage of electric power.
- Agents expressing local preferences and exchanging information with providers and other consumers can obtain a global view and can act optimally in both an individual and global sense.

5 References

1. S. Rahman. Integration of demand response with renewable energy for efficient power system operation. In IEEE PES ISGT Middle East Conference and Exhibition. 2011.
2. Iain S. Walker and Alan K. Meier. Residential thermostats: Comfort controls in California homes. Project Report LBNL-938E, Lawrence Berkeley National Laboratory, March 2008.
3. Jesse Berst. How cities (and their utilities) are blowing it: New research reveals three failures. Smart Grid News, December 2011. `http://www.smartgridnews.com`.
4. Sean P. Meyn, Matias Negrete-Pincetic, Gui Wang, Anupama Kowli, and Ehsan Shafieepoorfard. The value of volatile resources in electricity markets. In Proc. 49th IEEE Conference on Decision and Control (CDC), pp. 1029–1036, Atlanta, December 2010. IEEE.

5. Jeff Howe. The rise of crowdsourcing. Wired, 14(6), June 2006. http://www.wired.com/wired/archive/14.06/crowds.html.
6. Saifur Rahman. Integration of demand response with renewable energy for efficient power system operation. In IEEE PES ISGT Middle East Conf. and Exhibition. 2011.
7. Damien Challet, Matteo Marsili, and Yi-Cheng Zhang. Minority Games: Interacting Agents in Financial Markets. Oxford University Press, Oxford, UK, 2005.
8. Stanislao Gualdi, Matus Medo, Yi-Cheng Zhang. Crowd Avoidance and Diversity in Socio-Economic Systems and Recommendation. January 2013. arXiv:1301.1887
9. Vlad Tarko. Minority Games, 2012. http://news.softpedia.com/news/Minority-Games-38625.shtml
10. Karthik Iyer and Michael N. Huhns. Multiagent negotiation for fair and unbiased resource allocation. In Proceedings of the OTM Confederated International Conferences, CoopIS, DOA, and ODBASE, LNCS, pages 453–465, Washington, DC, October 2005. Springer.
11. Karthik Iyer and Michael N. Huhns. A procedure for the allocation of two-dimensional resources in a multiagent system. International Journal of Cooperative Information Systems, 18(3–4):381–422, September 2009.
12. FERC. National action plan on demand response. Docket AD09-10, Federal Energy Regulatory Commission, Washington, DC, June 2010. http://www.ferc.gov/legal/staff-reports/06-17-10-demand-response.pdf.
13. ESP. Energy service providers, 2012. http://www.csd.ca.gov/Programs/Energy
14. Johanna L. Mathieu, Duncan S. Callaway, and Sila Kiliccote. Variability in automated responses of commercial buildings and industrial facilities to dynamic electricity prices. Energy and Buildings, 43(12):3322–3330, December 2011.
15. Shengnan Shao, Manisa Pipattanasomporn, and Saifur Rahman. Demand response as a load shaping tool in an intelligent grid with electric vehicles. IEEE Transactions on Smart Grid, 2(4):624–631, December 2011.
16. Shengnan Shao, Manisa Pipattanasomporn, and Saifur Rahman. Grid integration of electric vehicles and demand response with customer choice. IEEE Transactions on Smart Grid, 3(1):543–550, March 2012.
17. Jay Alabaster. Japan group to build smart power grids that treat energy like network data, December 2011. http://www.computerworld.com/s/article/9222580
18. Mesut E. Baran and Ismail M. El-Markabi. A multiagent-based dispatching scheme for distributed generators for voltage support on distribution feeders. IEEE Transactions on Power Systems, 22(1):52–59, February 2007.
19. Luis Hernández, Carlos Baladrón Zorita, Javier Aguiar, Belén Carro, Antonio Sánchez-Esguevillas, Jaime Lloret, David Chinarro, Jorge J. Gómez-Sanz, and Diane Cook. A multi-agent system architecture for smart grid management and forecasting of energy demand in virtual power plants. IEEE Communications Magazine, 51(1):106–113, January 2013.
20. Manisa Pipattanasomporn, Hassan Feroze, and Saifur Rahman. Multi-agent systems in a distributed smart grid: Design and implementation. In IEEE PES Power Systems Conference and Exposition (PSCE), pp.1–8, March 2009.
21. Manisa Pipattanasomporn, Murat Kuzlu, and Saifur Rahman. Demand response implementation in a home area network: A conceptual hardware architecture. In Procedings of the 2nd IEEE PES Conference on Innovative Smart Grid Technologies (ISGT), pages 1–8, Anaheim, California, January 2012. IEEE Power & Energy Society.
22. J. S. Rosenschein and G. Zlotkin. Rules of Encounter. MIT Press, Cambridge, MA, 1994.
23. Steven J. Brams and Alan D. Taylor. Fair Division: From Cake-Cutting to Dispute Resolution. Cambridge University Press, Cambridge, UK, 1996.